

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 April 2002 (25.04.2002)

PCT

(10) International Publication Number
WO 02/33895 A2

(51) International Patent Classification⁷: **H04L 12/00**

(21) International Application Number: PCT/US01/32319

(22) International Filing Date: 17 October 2001 (17.10.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/241,450	17 October 2000 (17.10.2000)	US
60/275,206	12 March 2001 (12.03.2001)	US
09/903,423	10 July 2001 (10.07.2001)	US
09/903,441	10 July 2001 (10.07.2001)	US
09/923,924	6 August 2001 (06.08.2001)	US
09/960,623	20 September 2001 (20.09.2001)	US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US	60/241,450 (CIP)
Filed on	17 October 2000 (17.10.2000)
US	60/275,206 (CIP)
Filed on	12 March 2001 (12.03.2001)
US	09/903,441 (CIP)
Filed on	10 July 2001 (10.07.2001)
US	09/903,423 (CIP)
Filed on	10 July 2001 (10.07.2001)
US	09/923,924 (CIP)
Filed on	6 August 2001 (06.08.2001)
US	09/960,623 (CIP)
Filed on	20 September 2001 (20.09.2001)

(71) Applicant (for all designated States except US): **ROUTE-SCIENCE TECHNOLOGIES, INC.** [US/US]; 167 2nd Avenue, San Mateo, CA 94401 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **LLOYD, Michael,**

A. [US/US]; 160 Arundel Road, San Carlos, CA 94070 (US). **FINN, Sean, P.** [US/US]; 1533 Escondido Way, Belmont, CA 94002 (US). **KARAM, Mansour, J.** [LB/US]; 707 Continental Circle, #421, Mountain View, CA 94040 (US). **BALDONADO, Omar, C.** [US/US]; 700 Alester Avenue, Palo Alto, CA 94303 (US). **MADAN, Herbert, S.** [US/US]; 347 Blackfield Drive, Tiburon, CA 94920 (US). **MCGUIRE, James, G.** [US/US]; 2312 Gough Street, San Francisco, CA 94019 (US).

(74) Agent: **SUZUE, Kenta;** Wilson Sonsini Goodrich & Rosati, 650 Page Mill Road, Palo Alto, CA 94304-1050 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR THE ASSESSMENT AND OPTIMIZATION OF NETWORK TRAFFIC

(57) Abstract: A system for the assessment of network performance criteria, and applying this criteria to the classification of network addresses into appropriate ranges, using these ranges to consolidate performance measurements for the associated addresses, and applying these metrics toward the optimization of the network towards performance or policy objectives.

WO 02/33895 A2

METHOD AND APPARATUS FOR THE ASSESSMENT AND OPTIMIZATION OF NETWORK TRAFFIC

BACKGROUND OF INVENTION

5 Field of the Invention

This invention relates to the field of networking. In particular, this invention relates to the optimization of network traffic.

Description of the Related Art

10 The performance of packet-based networks (such as the Internet) is typically constrained by the various path-selection processes in place.

Manually configured forwarding tables are not flexible enough to compensate for variations in network performance. In addition, this approach represents a significant impact on scarce skilled resources.

15 Dynamic Routing Protocols are used to control the vast majority of forwarding decisions on large packet-based networks. These protocols are fairly adaptable to routing failures, but do not effectively compensate for variations in network performance.

Assessment and Optimization of network traffic is also subjective:
20 performance characteristics toward a given destination vary based on the location from which measurements are made. Each locale must identify their own local performance characteristics.

As the number of valid destination addresses increase, the feasibility of measuring performance characteristics toward all destinations in a persistent
25 and ongoing manner decreases to nil.

Consolidating the set of all individual destination addresses into sets, or ranges of addresses significantly reduces the number of tests that must be made in a given time interval. However, in many cases, even this optimization is insufficient: the current Internet routing tables include over 100,000 different
30 "routes", or address range identifiers.

The performance aspects of network paths to any given address will often have similar characteristics of the network paths to a nearby address. However, the task of identifying address blocks which are acceptably similar in

performance characteristic is subtle: network address ranges are typically not directly representative of the underlying network topology.

In most locations, the set of actively used address ranges is a relatively small percentage of the total address ranges available. This set of “active” routes is dynamic, and while technology does exist that can identify the set of active flows, in general these systems are constrained by both significant delays in reporting activity, as well as lack of integration into path selection systems.

Most network routing protocols in production use in the Internet are based on relatively primitive topological metrics such as “hop count” or “link costs”. Network performance metrics can include a wide variety of units and scales, which typically do not conveniently map to appropriate routing metrics.

SUMMARY OF THE INVENTION

A method and an apparatus for the optimization of internetwork traffic, including the selection of flows for optimization, quantifying appropriate network performance characteristics of these flows, calculating a performance metric from these characteristics, and updating routing or forwarding tables based on this performance metric.

A method and an apparatus for the optimization of internetwork traffic, including the selection of flows for optimization, quantifying appropriate network performance characteristics of these flows, calculating a performance metric from these characteristics, and updating a name resolution system based on this performance metric.

A method and an apparatus for the optimization of internetwork traffic, including the selection of flows for optimization, quantifying appropriate network performance characteristics of these flows, calculating a performance metric from these characteristics, and redirecting traffic flows based on this performance metric.

A method of classifying network addresses into appropriate Groups, based on evaluation of a selection of performance, policy, topology, and related criteria. Having defined said Groups, instantiate them in a server or other networking device, such that performance scores for observed traffic can be consolidated within the Group definition.

Additional features and benefits of the present invention will become apparent from the detailed description, figures and claims set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows an embodiment of a high-level overview of a specific embodiment of a network device deployed in an internetwork that is measuring performance characteristics of a subset of flows, calculating metrics based on those measurements, and changing the routing. Examples of how to change the routing include using a dynamic routing protocol (either an EGP like BGP or an IGP like OSPF) or reconfiguring the router, e.g., using route maps or static route statements. Note that the network device can be inside an existing router.

Figure 2 shows an example of needing to determine an exit point from an internetwork and some ways to compute the exit point.

Figure 3 shows examples showing network devices affecting routing across multiple internetworks and affecting routing inside a single internetwork.

Figure 4 shows some possible embodiments of devices that are communicating with each other, for example sending and receiving measurement packets.

Figure 5 shows one specific detailed embodiment of two devices, where each device is sending and receiving measurement packets as well as selecting a subset of paths.

Figure 6 shows an embodiment with more than two devices that are sending and receiving measurement packets to obtain measurements of performance characteristics of paths and to communicate measurements statistics about those paths.

DETAILED DESCRIPTION

In this document, we describe mechanisms that can be used to collect performance information regarding a set of flows, and re-route these flows appropriately. The flows may be between points within a single administrative domain, or at least one of the end points may be in some other domain. The mechanisms described here assume that it is desirable to observe the performance characteristics of the various paths over which it is possible to route the flows. This can occur in the case of organizations making use of network fabric they do not directly control; this includes the exchange of information with a remote end point across the Internet, or the exchange of information between end points belonging to a single organization over fabric owned and operated by a different organization (a Service Provider, or SP). These examples do not limit the application of the invention.

This invention includes mechanisms that implement some combination of the following parts:

1. Determine which flows need to be monitored
2. Collect performance information about these flows
3. Divert traffic, if necessary, by way of Edge Routers connected to multiple Service Providers
4. Route the traffic optimally to the end point
 - a. Choose an egress PoP where traffic can be forwarded
 - i. Definition of Equivalence Classes
 - ii. Centralized management of ECs
 - iii. PoP to EC Monitoring
 - iv. Winner PoP determination
 - b. Choose the best service provider to reach this PoP
 - i. Monitoring
 - c. Two-tier routing versus Dogleg (combo) routing
 - i. Routing Information Communication Protocol: BGP
 - ii. Choice of measures, and metric construction
5. Optimized Routing in the forward direction

These mechanisms can be implemented in many ways; examples include a single device that implements all of the required functionality. Other examples include a number of different devices, each being responsible for one or more functions. These devices can be located in the network of the domain that generates the traffic which we intend to measure (in this document, we call this domain the Content Provider - CP). Typically, such devices will be located in the CP's DMZ; alternatively, these mechanisms can be implemented in a Colocation Facility (CoLo). Each location is characterized by a different set of requirements. For example, in a CoLo, direct connections can be made between the collection devices and the border routers to which connect the organization whose traffic we measure. We describe implementation examples that correspond to both these environments.

In the context of one specific embodiment of this invention, we describe the hardware needed to collect measurements and bend traffic appropriately. Whether the mechanisms are implemented at a DMZ or Colo, the existing network infrastructure in which these mechanisms are being implemented can either be an autonomous system, or an autonomous sub-system.

In all of the text below, a Surfer merely denotes any IP endpoint that we do not control, and to which we need to develop routes on demand.

Monitoring Exit Traffic

In this embodiment of the invention, a passive monitoring and reporting device is needed, either in the enterprise DMZ or in a collocation facility (CoLo) rack; let's call this device the Collector. The job of the collector is to maintain simple state on selected flows emerging from the CP, and report statistics on them. In some embodiments of this invention, there is an Ethernet link to instrument, in which case packet splitters can be used; to such a device a computer that runs the Collector software can be attached. The software needs to receive signals about addresses (and possibly ports) for which to monitor flows.

In some embodiments, the Collector will track sequence numbers (looking for repeats, which can indicate packet loss) and SYNACK to ACK

time (as an estimate of edge-to-end RTT). It is then able to report these numbers, or track thresholds and report violations.

5 In some embodiments, a flow-based data collector (such as RMON II or NetFlow) can be used to identify the flows to be monitored; such a device can also be used to get performance characteristics of the flows.

10 In some embodiments of this invention, the device can select flows on the basis of input other than performance, such as user input, or the cost of the link. That is, the user can specify the flows that need special treatment. Also, interface statistics can be obtained from network; using this input, quantities such as available bandwidth, cost of usage of the links, or link bandwidth can be computed. Those skilled in the art can identify other metrics based on which flow determination and treatment can be based. Flows can be diverted when a given threshold in such metrics (such as cost or available bandwidth) is detected.

15 In some embodiments, some of all of the techniques above can be combined to yield a metric that determines whether a flow should receive special treatment.

In some embodiments of this invention, the set of flows can be provided by an external agent. In some embodiments, this external agent can be human; in other
20 embodiments, these flows can be provided by the external agent according to some criteria.

In some embodiments, the criteria for flow selection can be based on the monitoring of a routing table.

25 In some embodiments, the device will select flows whose destination network addresses correspond to some selected ranges of network addresses in the table.

In some embodiments, the selected ranges of network addresses can correspond to those flows that are of highest importance to the customer.

30 In some embodiments, the criteria for flow importance can be the rate of incoming requests.

In other embodiments, the criteria for flow importance can be the revenue that is generated as a result of the activity of the flow.

Those skilled in the art can identify other factors that can influence the importance of a flow.

In other embodiments of this invention, flows are identified by matching criteria defined on any portion of an HTTP request. In one implementation of such an embodiment, the well-known technique of identifying individuals or groups using cookies can be used.

In other embodiments of this invention, all flows can be monitored and/or diverted. That is, there is no requirement that some part of the flow remains non-monitored and/or non-diverted.

The Collector communicates with another processes, referred to below as the Bender, that is responsible to divert traffic for which optimized routing is needed.

In one embodiment of this invention, the precise placement of the Collector depends on whether it is implemented in the CP's DMZ, or next to the CP's border routers in a CoLo. In some embodiments, the Collector can include passive, fail-safe monitoring hardware, and it may be placed in the flow for all the outbound traffic (such as within a DMZ). In other embodiments, it may be required that all traffic that needs to be monitored will be diverted towards the Collector. (See below.)

Bending A Flow

In some embodiment of this invention, the second device is referred to as the Bender, which can (a) receive input from the CP and the Collector, (b) collect stats on various routes to and through the different available service provider links, (c) make decisions about which flows require diversion, (d) obtain those flows from the CP, and/or (e) re-route them depending on the input received from the Collector. The Bender may, but need not be physically distinct from the Collector described above. In some embodiments, stage (d), for example, (which gives rise to the name Bender) is optional in some implementations. If the flow is to be modified in any way (Network Address Translation (NAT), insertion into a tunnel, redirection onto a reserved additional pipe), then diversion as described may be necessary when the existing network devices do not offer the features. However, if the surrounding network fabric

can readily be altered by the Bender to achieve the desired effect, step (d) can be present or omitted.

5 In some embodiments of this invention, all flows monitored by the collector can be sent to the bender. In other embodiments, only a sub-set of these flows can be monitored. All techniques used above for the determination of which flows to monitor can be used to determine which flows need to be re-routed.

10 In some embodiments of this invention, this mechanism includes a network device (for example, some kind of a router), under control from some processing unit (a computer) that runs the appropriate software. In one embodiment, the computer is responsible for (a) through (c), and then communicates with the router to make (d) and (e) happen appropriately. In one embodiment, a human can intervene on the router to effect stages (d) and (e), on the basis of reports or signals generated by the computer performing steps (a) through (c). Those skilled in the art will readily identify different packaging schemes for this invention. In this embodiment of the invention, we assume that a given flow is identified by its destination IP address; those skilled in the art can identify other ways to identify flows. When the CP wants to signal that traffic towards a particular IP address needs to be optimized (see below), it should communicate directly with the Bender. The Bender can communicate about qualified flows with the border routers that connect to the different available service provider links. It may be necessary for the external fabric to notify the Bender of a flow, depending on the CP's particular choice of signaling method (e.g., an API, through, for instance, HTTP requests.)

20 25 The Bender will inform the Collector that a particular flow is qualified, and the Collector will begin to monitor the flow. The Collector returns statistics to the Bender. In some embodiments of this invention, these statistics may encompass more flows than are strictly qualified; indeed, it is possible to infer performance problems in non-qualified flows, and use these proactively. The Bender will then implement logic that determines whether the qualified flow needs to be diverted. In a simplified embodiment of this invention, all qualified flows can be diverted to the Bender. This would significantly simplify the function of the Collector.

The Content Provider's Choice of Flows

A number of possible methods can be used by the Content Provider to inform the collector that a flow is qualified for optimization. Various embodiments have one or more options for the Content Provider, including:

- 5 1. No notification – all active sessions are qualified Surfer addresses
2. Static pre-configuration of qualified addresses
3. URL adjustment, where a Web Surfer will attempt to access the next page via a special address, and we use that as the signal that they are qualified
- 10 4. Web page embedding, where the Surfer will download a graphic or other special object, and the access is taken as the qualification signal
5. Calling specialized code within the CP's Server applications which communicates the address to the Bender
6. Calling a specialized method on one of the Bender devices
- 15 7. A remote call, in the form of an HTML post operation.

Those skilled in the art can identify other methods used by the Content Provider to inform the Collector that a flow is qualified.

Some embodiments implement Option 7. A simple web interface is offered as part of the Bender. Additional logic will take care of performing post operations on the web transactions. A secure web form interface can also be offered as part of other Options for flow qualification; this allows the Content Provider to either experiment with the service, or intervene to control it once it is running. In a simplified embodiment of this invention, the HTML post interface supports one operation, namely to deem one IP address a qualified address for a specified period of time. This period of time is specified by the Content Provider, which enters for this purpose a TTL (Time To Live) for the qualification; it may be desirable to keep the TTL value small; it may also be desirable to provide a default, possibly at around 5 minutes.

30 In another embodiment of this invention, a more complex interface can be built, that includes the ability to withdraw qualification, and/or to specify additional parameters such as port numbers, relative priority versus other qualified flows, and/or a general statement of QoS requirements.

Deciding to Divert

In this embodiment of the invention, the Collector sits in the flow of packets, but is strictly passive, and can be built in such a way that even power failure does not compromise the monitored network fabric. (Line splitters and Switch span ports are two examples.) In some embodiments of the invention, the Bender sits out of the normal flow direction, and so must attract any qualified flows.

In some embodiments of the invention, forcing a qualified flow to cross the Bender requires a routing protocol between the Bender and the nearest Content Provider routers. An implementation involves the injection of /32 routes (or an aggregation thereof) for all qualified flows that need to be bent. In some embodiments of this invention, precaution can be taken to withdraw or age out individual advertisements, to avoid the excessive growth of routing tables in the Content Provider routers. One implementation is to have the processing unit cause the router to advertise the qualified destination; one way of implementing this effect is to have the processing unit reconfigure the router (e.g., via telnet) to add a static route. The router may be previously configured to redistribute such statics into the agreed routing protocol running between the Content Provider's edge routers. Many routing protocols can be used in this context. It may be preferable not to be in the normal IGP of the Content Provider, so as to minimize the distance the potentially rapidly changing routes propagate. In some embodiments of this invention, a common protocol with rapid convergence properties (e.g., OSPF) is ideal.

Adding the static route and advertising it will cause the qualified flow to divert to the Bender (after a route convergence time). In case the IP address is used to determine qualified traffic, and if this IP address is shared between qualified and non-qualified traffic, unqualified traffic may also be included (e.g., because of Port Address Translation - PAT). In other embodiments of this invention, qualified and unqualified flows within a single IP address can be separated.

Forwarding Qualified Flows Through Optimized Service Provider Routes

Once a flow has been diverted, it can be forwarded through an optimized route.

The static used to divert the flow is the first piece of tunable routing. In some embodiments of the invention, the router component of the Bender has connectivity to edge routers which are in turn connected to various service provider links, to which qualified traffic will be sent. (In this document, we will refer to this set of routers and links at the edge of the domain as a point of presence – PoP.) In this document, we focus on an implementation that assumes the existence of GRE tunnels between any connected PoPs. Those skilled in the art can identify other implementations that do not involve the use of GRE tunnels; for example, connected sites may have IPSec fabric between them, or no tunneling mechanism at all. In the case of GRE, static routes used to divert traffic can point directly into one (or more than one) of these GRE tunnels, based on a decision process in the processing unit of the Bender.

Recursive Static Routing

In other embodiments of this invention, another layer of indirection can be added; one such implementation involves the use of statics to remote IP next hops. The idea is to decouple the static that bends the traffic towards specialized devices from the details of how it is to be forwarded. Such recursive statics provide a useful simplification for minor performance cost.

The specific implementation details of these recursive static routes can depend on whether the Bender devices are inside the Content Provider's DMZ, or in a neighboring CoLo rack. In the first case, it is sometimes desirable not to move a full feed of performance data in to the Bender, which could prevent the implementation of complex edge-to-edge routing decisions. When in a CoLo facility, the Bending devices are presumably fully connected to the PoPs, rendering the implementation of full routing decisions to be an easier task (as described below under routing for core PoPs).

In the DMZ case, some embodiments of this invention involve having a process that adds static routes to the router (in addition to those static routes

advertised to divert the flows to the Bender, described above). These processes may involve the addition of several specific statics, with next hops that spread across the GRE tunnels leading to various PoPs. In some embodiments of the invention, this indirection is outside the GRE tunnels themselves, hence
5 preserving the property that a single flow will all go to the same next PoP, even if load sharing across tunnels per flow is in use.

In a CoLo example, the process for selecting nearby PoPs is, in some embodiments of this invention, unnecessary, and the initial statistics used to attract flows can simply point to the core-facing router in the PoP. The core-
10 facing router will already have selected a high-performance route to the Surfer, as described below.

The Core Fabric for Optimized Routing

Whether the interface with the Content Provider is within their DMZ, or out in a CoLo, packets will eventually arrive at a PoP connected to a core
15 Internet fabric, where routing optimization can take place. In some embodiments of the invention, the problem can be decomposed into three pieces:

1. Choosing one or more candidate egress PoPs, and
2. Choosing routes to those PoPs.
- 20 3. Implementing routes to these PoPs

In general, Choices 1 and 2 are not orthogonal decisions, since the egress can depend on the cross-core performance. In some embodiments, methods for steps 3 can include automatic changes to an IGP, to BGP, or manual changes to either. In some embodiments, these changes are in the form
25 of router configurations that statically alter the routing of flows (methods including use of route maps and/or use of static route statements). In other embodiments, changes are in the form of router configurations that adjust the processing of dynamic routing updates. In some other embodiments, changes to routing are determined within the forwarding device, and are applied directly to
30 its routing table.

Choosing one or more candidate egress PoPs

When routing towards general Surfer addresses, it is, in general, undesirable to sustain individual routes to every Surfer address, since this can cause significant growth of Internet routing tables. Hence, in some embodiments of this invention, the implementation is constrained to be designed for scale. In general, many neighboring hosts can be routed in the same way. However, in some other cases, load sharing may be used, so it will not be the case that any two flows to the same subnet will follow the same actual path across the core.

10 Address “Equivalence Classes” (or “Group”) Definition and Applications

Some embodiments of this invention use the notion of “Equivalence Classes” (ECs), also known as “Groups”.

An “EC”, or “Group” is a set of network addresses which are to be treated equally by the system for the purposes of performance reporting, and/or performance optimization.

The set of network addresses within a group may include any combination of individual network addresses, and ranges of contiguous network addresses.

In some embodiments of this invention, the network addresses in question will be Internet Protocol (IP) addresses. In this case, the traditional “a.b.c.d/m” notation for Internet routes can be used, facilitating communication with conventional Internet routers.

The assessment of a set of network addresses to identify appropriate Group membership requires the identification of appropriate criteria of each network address to evaluate.

In general, address classification criteria falls in four (4) distinct classes of attributes:

1. Performance-related
2. Administrative Policy
3. Topological Proximity
4. Pre-Existing Set Definitions

Performance-related criteria include those parameters which may be observed of traffic to and/or from an address or address range. Examples of such parameters include: delay, jitter, loss, retransmission, and network load.

These parameters may be observed as a result of active measurement (transmission of test traffic); or passive measurement (observation of existing flows).

The range of Administrative Policy criteria encompasses a range of considerations that may, or may not be directly observable within the network fabric itself. Examples include:

Cost Considerations: a specific address range, or set of address ranges, may benefit from reduced costs from a specific provider. Examples of this case include directly negotiated "Private Peering" agreements.

Address Aggregation Policy: In some cases, scalability considerations of individual networking devices, and the system as a whole, result in desires to limit the number of address ranges to be managed. In order to keep the number of address ranges reasonable, defining "minimum set criteria" for address ranges is sometimes required. Examples of this case include "minimum set size of 255 addresses"; which is to say, addresses, and address ranges involving less than 255 addresses, will be aggregated into a single set. In some cases, Address Aggregation policy varies across the range of managed addresses. For convenience in identifying sets, Address Aggregation policy often is limited to ranges of contiguous addresses.

Administrative Policy may also encompass user-definable criteria.

Topological Proximity criteria allows network addresses which are either topologically contiguous/adjacent, or topologically dependant, to be considered in the definition of Groups. Topological Proximity can be empirically observed through active testing of the network fabric, or inspection of appropriate network documentation. In addition, Topological Proximity can in some cases be inferred from the relative contiguousness of addresses.

Pre-Existing Address Range definitions may be usefull in the Group definition process: the use of such definitions, if available, may increase the efficiency of bootstrapping the system, and may also allow the system to integrated more effectively with internal or external routing or forwarding devices. In some cases, Pre-Existing Address Range Definitions may map directly (1:1) to Groups; in other cases, a Group may include more than one of the existing Address Definitions. In some examples of this case, a source of routing table information may be identified as the reference source of Group Definition information. In some emodiments of this invention, Group Definitions are kept synchronized with a dynamic Reference Routing Table.

In some embodiments of this invention, it may be desirable to restrict candidate Reference Routing Table entries to be considered for Group membership to contiguous address blocks.

In some embodiments of this invention, combining discontiguous entries in the Reference Routing Table into a single Group membership is acceptable.

Those skilled in the art can identify other metrics which may be considered in the definitions of Groups.

In some embodiments, some or all of the attributes above can be combined into a composite metric to be used for group membership evaluation.

Individual network addresses, or network address ranges are evaluated in the context of the selected criteria, and placed in a Group based on this evaluation.

In some embodiments, the definition of groups may at least partially computed with the help of a "clustering" algorithm, which includes one or more of the classification criteria described above.

Updates

Assignment of network addresses, and network address ranges, into appropriate Groups may need to be reviewed and repeated based on changing network conditions including advertised address ranges, network performance characteristics, and user-defined policies. This process manually initiated, or it may be triggered by the detection of changes in network conditions; or may be scheduled to occur on a regular basis.

Central vs. Distributed Definition of Group Membership

In some cases, it may be important for all PoPs to agree on the specific definitions of each Group. In addition, the computation required for flexibly
5 adjusting the Group sizes is both CPU and data intensive. Consequently, in some embodiments of this invention, the management of the Group list is centralized. One implementation involves:

1. A central server
2. A data warehouse for historic performance data
- 10 4. A communication system that can inform all PoPs of changes in Equivalence Classes as they are found.

One implementation involves allowing the centralized Group manager to determine which PoP advertises a new Group Definition first, based on its
15 complete information base. This implementation will allow PoPs to coordinate, and will hence prevent them going out of sync.

In other embodiments of this invention, Group Management and definitions can be independent in different POPs. That is, the actions of creation, modification,
20 and deletion of the groups are performed at sites distributed across the different PoPs. One simple implementation of such an embodiment can involve a PoP advertising to all other peers in the network whenever it makes changes to the EC list (that is, upon the creation, modification, and deletion of a group or a set of groups).

25
Specific embodiments of this invention will differ in the details of their Group Definitions, depending on the range of IP addresses that are deemed "equivalent".

Application of Groups to Routing Tables

30 In some embodiments of this invention, it may be desirable to enforce a direct 1:1 mapping between Group Definition and a Reference Routing Table.

In some embodiments of this invention, the Reference Routing Table may be obtained from within the local Autonomous System.

In some embodiments of this invention, the Reference Routing Table may be obtained from outside the local Autonomous System.

5 In some embodiments of this invention, it may be desirable to allow Groups to refer to exactly one or more entries of the Reference Routing Table.

In some embodiments of this invention, it may be desirable to allow Groups to refer to one or of the entries in the Reference Routing Table, in addition to one or more individual addresses, each of which may, or may not, be
10 explicitly included in the Reference Routing Table.

In some embodiments of this invention, it may be desirable to allow Groups to refer to address ranges which were are not present in the Reference Routing Table. In some cases, this may involve address ranges within Group definitions which are broader than those currently present in the reference
15 routing table.

In some embodiments of this invention, it may be appropriate to constrain the address ranges in Group definitions such that they refer to ranges which are no larger than the existing entries in the Reference Routing Table.
20

Monitoring Equivalence Classes in Possible Egress PoPs

Given the set of Equivalence Classes into which the Internet is divided, this section describes how each Equivalence Class is routed. As a first part of this process, in embodiments of this invention, PoPs measure their distances to
25 an EC via various external routes; that is, if PoP A had packets heading to EC B, and it chose to drop those packets out to the Internet, how well would they perform?

Embodiments of this invention can implement one or more of the following monitoring options:

- 30
1. A direct attempt to ping the Surfer address.
 2. Traceroute towards the Surfer.

3. Coercing the Surfer to communicate with a PoP as part of the conversation, for example by embedding a particular GIF as part of their Web page. (Telnet or other protocols could also be used.)
4. Asking the Surfer to run some code that initiates connections to candidate PoPs. (The motivation for the Surfer would be the improved performance if they were to do this. Note also that only one or a few Surfers within an EC need to accept the code to get a measurement for the whole EC.)
5. Passive monitoring of the flow, either within the Ingress PoP or elsewhere. (Note that in some implementations the Ingress PoP will need to be able to cope with only seeing unidirectional communication, since we cannot always guarantee that return packets will go via the same PoP.)
6. The simple observation of the Surfer's routed choice of Ingress into the Optimized Routing Infrastructure. (Clearly, this and the passive monitoring approach require that the forward flow actually impinge upon the core fabric.)
7. An estimate of geographic distance to the Surfer, in some embodiments obtained from databases which correlate Internet addresses with locations
8. The regular BGP answer for distance to an address range. Note that this has the useful property of always generating an answer if no other basis can be found. (This will require a little adjustment once our EC's no longer match existing BGP routes, but the task is not too difficult.)

Those skilled in the art will identify other monitoring techniques.

In some embodiments of this invention, it may be desirable to ensure the results of these tests are comparable between PoPs.

The implementation described in this document includes a PoP assessing its "distance" to a given EC autonomously along every Egress connection it has to the Internet, and to summarize the result as a simple metric for communication to other PoPs. That is, in this particular implementation, other

PoPs are not fed individual data points or estimates of network quantities such as delay. In some embodiments, this detailed data for each individual Surfer address is fed to the centralized EC Manager, since it is the basis under which EC's can be separated or joined. However, this exchange is distinct from PoP to PoP communication, since routing does not need the fine details; a simple summary of a PoP's "goodness" for a given EC is all that is required.

In some embodiments, measurements can be boiled down into a single statistic. In some cases, it may be important to rank the measurements by their expected quality. In one example implementation, the quality order is:

1. Direct ping measurements to the Surfer
2. Spoofing races; that is, a technique where multiple PoPs respond to a query, all using the same address ("spoofing"). For some common network protocols (including at least DNS, HTTP and TCP), when the answer to the query is used, it is possible to tell which PoP's answer arrived first, and hence is the closest.
3. Observed Ingress PoP
4. Passive monitoring at Ingress PoP
5. Traceroute
6. BGP

Given this hierarchy of expected quality, a metric can be constructed, that can be propagated in existing protocols by "concatenating" the measurements; that is, separate measurements can occupy different bits counting from the MSB in an integral metric value. This particular example implementation does not require that the concatenation be strict; overlap of values is allowed, since there is no need to reconstruct precise and separate estimates when these measures are received. In this implementation and other embodiments, the goal is to find those PoPs that are better than others, which are preferred. Some embodiments can consider this property to hold even if the measurements are heterogeneous in method or time.

In the implementation described in this document, low resolution in the metric may be considered acceptable, since this tends to diminish the likelihood of reacting to small changes in performance. The goal of some embodiments is to find several routes, none of which are appreciably sub-optimal, and any one

of which can be pulled out of service rapidly (with the others taking over) when sudden performance degradations such as packet drop storms are observed. In those embodiments, paths whose measurements vary by insignificant amounts are treated as equal; in some implementations, this can be achieved with a composite metric of low granularity.

When to Monitor

In some embodiments, it may not be feasible for all PoPs to monitor all EC's across all exit paths at all times. For such embodiments, a decision scheme is needed to trigger the previously described statistics. A few selection techniques can be combined, including at least the following:

1. When a new Surfer address is observed in the PoP (no matter what role the PoP has, since any PoP can choose to abandon packets to un-optimized forwarding), its EC is measured (preferably using the Surfer's actual address in applicable tests)
2. BGP metrics indicate the EC should be "close"
3. Random sampling

Those skilled in the art will identify other such techniques.

In some cases, combining all of these techniques can still generate too much test traffic. Hence, in some embodiments, weighted sampling can be implemented; in one such implementation, EC's would be ranked by their priority for testing, and tests would concentrate on the top few at any given moment. In this implementation, initial weighting can be derived from BGP: observation of a flow crossing the PoP causes a large increase in priority (and a record is made of the individual address within the EC). In different embodiments, the weighting can either be cumulative (so that repeat flows from a given EC increase the testing focus applied) or old observations eventually time out.

In other embodiments, more sophisticated feedback can also be added to the EC weighting scheme, using, for example, the data a PoP receives from other PoPs. In one such implementation, a PoP that has at least one measurement for a given EC, compares that measurement to those coming from other PoPs; if its own measurement appears to be close, then it intensifies

monitoring; if it is far from optimal by a large margin, then the EC drops far down the priority list. In some embodiments, the EC Manager can suggest EC's for a PoP to monitor, based on the warehoused performance data it is analyzing.

5 *Core PoP to PoP Monitoring*

In this case, both end points of a cross-core links are under our control. In some embodiments, it is desirable to generate cross-core measurements that are comparable with the Egress PoP measurements described above, so that path comparisons can be made. In some embodiments of this invention, an overlay network can constitute the network infrastructure that links these PoPs to each other. In the following implementation, we assume the use of GRE tunnels across POPs. In some embodiments, this overlay network can be set up as a single Autonomous System. In this implementation, the overlay network is set up to be single BGP Autonomous system. On the core, different quantities of interest along all direct paths between all PoPs can be directly monitored. Possible measurements cover, and are not restricted to one or more of the following:

1. One-way delay
2. One-way packet drop rate
- 20 3. One-way jitter
4. Available access bandwidth. (In some embodiments, this quantity can be estimated via the output buffer fullness on the physical interface at the sending end of the inter-PoP tunnel.)
5. Complete tunnel failure

25 Other embodiments can use two-way measurements instead of one-way measurements. These two-way measurements could include, but are not restricted to delay and packet drop. Those skilled in the art can identify measurement tools readily available for this purpose. In some embodiments, one-way measurements require clock synchronization between end-points; in some embodiments, GPS feeds can be used in the various PoPs. In other
30 embodiments, an NTP clock feed of sufficient quality can be obtained into each PoP.

In some implementations, it could prove beneficial to arrange these results in a low-resolution metric, where effectively similar choices come out with identical numbers. In other instances, the measurement precision will be required to accurately measure the effect on end-user application performance. (Different application measures differ in their significance to end-user application performance; also, various applications are relatively dependent on some quantities, but insensitive to others. For instance, web traffic is immune to jitter until it is severe enough to cause packet reordering, but is very sensitive to packet drops. Conversely, voice and video are relatively tolerant of packet drops, but are extremely sensitive to jitter.)

In some embodiments of this invention, the routing fabric is split with different types of traffic (e.g., jitter sensitive and drop-sensitive types of traffic) utilizing separately optimized routes. One implementation signals which method of handling is best for a given application through the ToS bits in the IP header. In this implementation, the CP, or a process at the CP's site could mark a special ToS value indicating the optimization required for the flow. In some embodiments, the ToS bits can be set flow by flow, using, for example specific knowledge of how the CP can separate application traffic (by port number, IP address, egress point, farm location, etc). However, it may frequently be the case that a CP will only have one class of traffic to put into the routing fabric, so in some embodiments, a constant marking rule can be applied. In all these cases, the traffic can be placed into the right routing system based on the CP-marked sensitivity type, using for this purpose a relatively straightforward use of Policy Based Routing (PBR) at the first PoP.

In other embodiments, the measures of different network parameters (e.g., jitter and loss) are mixed. In such an implementation, the mixture of different measures could hunt for physically short paths that are under-utilized (due to inefficiencies in conventional routing); a relatively under-used route shows improvements for both drop-sensitive and jitter-sensitive traffic. This architecture is viable since experience has shown that high jitter and high loss can be highly correlated phenomena. (That is, the various network degradations are frequently correlated. Real observations show that the correlation is not

ironclad, but appears to be present.) For example, voice flows are known not to be particularly sensitive to drops, yet they are sensitive to many of the network conditions which induce drops.)

5 One other item regarding the metrics deals with their ordering. In some embodiments, Jitter and Delay can be merged.

In such implementations, the proposed order for components of the metric becomes:

1. Packet drops
- 10 2. Delay + Jitter
3. Bandwidth

15 In some embodiments, Failure can be omitted, because it does not cause a metric change, but forces a rapid route removal instead.) In some embodiments, each remote PoP generates only a single metric value from the measurements, even though several direct paths are measured. In some
embodiments, a numbers across paths could be combined. In this document, we described an implementation for which the metric from the single best path will be used.

20 Combining these measures with the Egress metric is discussed in the next section.

In simplified embodiments of this invention, cross-core routing and choice of Egress PoP (using EC measurements) can be done independently. In such embodiments, PoPs are based on their proximity to the Surfer, not on their
25 optimality when the flow is at another given point in the network. In other implementations, cross-core routing and choice of Egress PoP can be combined within the same routing infrastructure.

Propagating Route Decisions Between PoPs

30 Now that the performance of various PoPs to an EC, and the PoP-to-PoP characteristics have been measured, a system for communicating these findings between PoPs is needed, so that routing decisions can be made. In some

embodiments of this invention, BGP can be used. Indeed, BGP is a communications protocol that is known to be capable of carrying the size of load such an infrastructure would generate; it also provides a number of useful features into one package (publicly available specs, modifiable source code for
5 Unix implementations, and support for "third party" next hops).

In one such implementation, each PoP will be equipped with an implementation of BGP. These BGP sessions should communicate with each other; in some implementations; all PoPs can peer in a full mesh of IBGP peerings. In other
10 implementations where the number of PoPs is large so that a full mesh becomes impractical, BGP Route Reflection can be implemented. In such an infrastructure, each PoP can advertise routes to all EC's that it has monitored, giving an address within that PoP as the next hop. Most attributes in that BGP message can be left at default values; in some implementations, only the 32-bit
15 Local Preference field is modified as to communicate the measurements described above for the closeness of the PoP to the EC. Those skilled in the art can identify other fields that can be used for this purpose (e.g., weight, MED, etc.) In implementations that use Local Preference, desirable routes should have higher values; in some embodiments, a scheme based on an ideal number of
20 points from which "demerits" are subtracted, for example when packets are observed to drop, or delay crosses a series of thresholds, can be implemented.

In embodiments using BGP, the BGP process will also receive routes from all other PoPs describing their performance as an Egress for some EC's. In implementations of these embodiments that combine cross-core and edge
25 routing measures, the receiving process can modify an incoming route, adding to the Local Preference a value measuring the cross-core distance from the sending PoP. However, in embodiments where both routing processes act independently (as described above), Egress PoPs can be chosen purely for their good performance towards a Surfer, regardless of cross-core performance.

30 In such implementations that use BGP, the computer can learn over time about possible Egress PoPs for each EC, and conventional BGP cost comparison can efficiently sort through for one or several winners. In some implementations of BGP, ties will force tie breaking. In other implementations,

multiple PoPs deemed “optimal” simultaneously could share load, which will have the advantage of adding redundancy, useful in the event of sudden degradations.

5 In some embodiments, where the rate of routing updates is large, throttles can be added. One specific implementation could avoid unnecessary advertisements of minor changes in route metrics. One such method is to fail to report changes where an underlying measurement has changed by less than a given percentage of the original measurement (e.g., 5, 10, or 20 %).

10 In some embodiments, the BGP updates can be sent directly across the core fabric. In other embodiments, this data can be moved out of the way of real user traffic (since it may be bursty, and will certainly increase congestion in the most active PoPs). In some such implementations, a relatively low-bandwidth communication infrastructure can be used to allow all PoPs to communicate. Such implementations can use a full meshing across peers. Other
15 implementations can use a hub and spoke pattern into the EC Management site. In case high resilience is needed, backup routes over the Internet can be designed. That is, for periods when the EC Management site is unreachable, routing traffic can be mixed with user flows.

20 *The Master BGP Local Preference Metric*

The BGP Local Preference Metric must contain information about Egress PoP selection. In some implementations, it does not deal with specifics of any given path towards the Egress PoP, or any exit interface from that PoP. As previously described, the metric needs to factor in one or more of the
25 following metrics: Packet Drops, Delay + Jitter and Bandwidth. In some embodiments, the metric needs to support an addition operation, so that a summary cost can be made to combine the Egress PoP cost to get to the EC with one or more PoP-to-PoP paths. Also, the use of BGP Local Pref forces semantics where high numbers are better – that is, a route gets a “goodness”
30 score, and as the path gets longer or less good, the score must decrease.

In some embodiments, bandwidth is not explicitly signaled between PoPs. (Given the details of tunnel selection (discussed later), it is non-trivial to interpret a Bandwidth measure directly. Also, the quantity may vary widely

even when the network is behaving well, since we can only measure Bandwidth remaining on the links in and out of our PoPs, not along the tunnels themselves, and loading on these links will fluctuate, while Packet Drops, Delay and Jitter may prove relatively stable much of the time.)

5 In some embodiments, bandwidth consumption can be used as an early warning sign, indicating the onset of link saturation and packet drops. A PoP has a number of options at or near link saturation. Choices include, but are not restricted to:

- 10 1. Re-enable a disabled tunnel leading to the remote PoP, which uses a different (unsaturated) link. This link could presumably involve degradation in average performance, or else it would already be in use, but the degradation may be deemed mild compared to the needed gain in bandwidth.
- 15 2. Utilize a tunnel leading to a different PoP, again over a different physical link, so that the other PoP can forward more effectively towards the chosen Egress PoP.
- 20 3. Send the packets out to the Internet over the locally best exit peering, which is not suffering saturation (presuming the packets are bound for the Surfer, not the Content Provider).
- 20 4. Use an "escape hatch", being a peering to an ISP over which no tunnels are constructed. Such escape routes may not be present in an implementation, but if the overall system has any routing problems such as delayed convergence, looping, or hysteresis, this may prove a cost-effective means of controlling adverse conditions.

25 Those skilled in the art can identify other such choices. In some embodiments, it is necessary to proactively assess degradation, in which case such predictions are made directly by the local PoP and factored in to the advertised scores.

30 In some embodiments of this invention, the performance-based, application-specific metric derived above could be modified using other variables such as overall network usage, individual network link usage (absolute or relative) or link usage cost. For example, interface stats of network devices (e.g., netflow on routers) could be used to determine current load on all available links, from which an estimate of the cost can be computed. The

performance-based metric can then be penalized to take into account cost factors. In some embodiments of this invention, these cost parameters can be manually entered/tuned by the user. In particular, in some embodiments, the user can explicitly penalize or favor given routes.

5

Enforcing Two-Tier Routing

In this section, we describe a class of implementations of this invention, where cross-core routing, and choice of egress PoP are treated independently. (We refer to such an implementation as Two-Tier Routing.)

10

Once the BGP process at a given PoP has learned how to route a given EC, it next needs to enforce that choice on the routers in the traffic flow. In this implementation, BGP is used for this purpose: that is, the computer feeds in the routes via IBGP (implying that the router is a Route Reflector Client of the computer), preserving the next hop address that comes from the remote Egress PoP.

15

The router now has to forward a packet to a remote IP next hop; in this situation, it will perform route recursion, allowing load sharing and resilience. In some embodiments, the BGP communication allows PoPs to decide which other PoPs to use. In other embodiments, a more local process can be inserted to offer the best path towards the selected PoP. Recursion inherits the load sharing behavior of routers, where flows can be split over multiple routes, but always with $1/n$ of the flows going to each route. In some cases, a $1/n$ scheme is a viable initial offering. In other embodiments, differential weighting can be required, in which case it is possible to duplicate the router mechanisms to allow moderate diversion from $1/n$. In one specific implementation of this idea, multiple similar routes can be inserted to point to the same tunnel. Those skilled in the art can identify other ways of implementing differential weighting.

20

25

30

The process of injecting specific static routes towards remote PoPs derives its decisions from the same PoP-to-PoP measurement previously described. Since all paths were probed to generate the simple cross-core metric, these detailed measurements can inform the choice of tunnel after BGP has selected the egress PoP. In the implementation described here, the BGP process

that monitors inter-PoP tunnels can keep track of which tunnels are optimal, and whenever changes occur, can reconfigure static routes on the router in the packet flow to reflect the new information. These static routes have destinations to match the BGP next hop fields (the addresses provided originally by the remote PoPs in their advertisements). The statics' own next hops point directly to GRE tunnels, which offer transit to the far PoP. In some implementations, the reconfiguration of the statics can be done by a human on the basis of reports generated by a monitoring process at the PoPs.

In summary, routing towards the Surfer using this class of implementations of two-tier routing includes one or more of:

- Performance monitoring of EC's from candidate Egress PoPs
- Performance monitoring of PoP-to-PoP GRE tunnels
- Combination of these values and PoP-to-PoP route computation within BGP processes in each PoP
- Insertion of the BGP routes into the forwarding router
- A lower layer of static routes, given to the forwarding router by a local process, which has final authority about how a packet should reach the selected far PoP

Using Dogleg Paths (Combo Routing)

In this section, we describe a class of implementations of this invention, where cross-core routing and choice of egress PoP are combined in the same routing infrastructure. (We refer to such an implementation as Combo Routing.)

Some embodiments of this invention, described in this section, have the ability

to route between PoPs without restriction to direct tunnels. In some

embodiments, this is achieved within the lower level static routes, which

determine the actual forwarding decision to Bag's chosen Egress PoP. In some

implementations of such embodiments, the system that is monitoring the direct paths to other PoPs is extended to generate a moderate amount of test traffic

down alternate paths. In some embodiments, it is excessive to ping via every possible path, so another weighted random technique can be used; in one such implementation, testing can be limited to a fixed ration of samples per minute, weighted heavily to the historically best performing alternate routes, but

occasionally sampling poor or unknown performers in search of significant changes.

In some embodiments, statics will be added that appear to route to an Egress PoP, but in fact follow a tunnel to an intermediary; in such
5 embodiments, the intermediary PoP can follow its own local routing decision to forward the packets. If no additional loop suppression mechanism is used, such an implementation could conceivably develop loops; indeed, since the various PoPs are deliberately set up to make autonomous decisions, it is possible for the sender and the intermediary to disagree about the correct Egress PoP. In most
10 cases, this is not a problem, but in rare cases a routing loop might develop. In environments where the desire to keep the probability of generating loops as low as possible, some embodiments of this invention can use loop suppression algorithms. An implementation of such a loop suppression technique includes the following: the technique is to disallow dogleg paths of more than three PoPs
15 in the core. In such an implementation, the intermediary PoP must know that the packet has already been sent off the direct path once, and it must limit its choices to direct paths to the Egress. Similarly, the Egress PoP must know that the packet has followed a dogleg, and it must therefore release the packet to the Internet (or deliver it directly without visiting another core PoP).

20 In some embodiments, the necessary communication can be achieved by creating "floating" GRE tunnels in addition to the existing direct PoP-to-PoP tunnels. In a given PoP, there will be several tunnels leading directly across the fabric to any particular remote PoP; at least one extra tunnel is declared, to a
25 special address in the remote PoP.

 In some embodiments, when probing shows that a candidate way to a remote PoP goes via an intermediary, static should be inserted to temporarily force the floating GRE tunnel between the sender PoP and the remote PoP into a direct path to the intermediary. In such embodiments, the choice should be
30 specific; the next hop used for the BGP routes cannot be used, because this can reintroduce loops. Each PoP can be pre-configured with routes that direct anything from the special subnet for a PoP to the direct links to that PoP. In this

way, the sender PoP can establish a tunnel to the remote PoP that is coerced to cross precisely one other selected Intermediary PoP.

5 In such embodiments, the selection of the direct paths between sender and intermediary, and between intermediary and remote, can be simplified to statics that load share across all tunnels. In some embodiments, the router removes any statics to paths that have no connectivity; in other embodiments, they can make use of the dynamic performance information we are already collecting for optimal choice of these statics.

Flows in the Forward Direction

10 In this final section, we describe one aspect of this invention, which involves the control of flows from the Surfer (i.e., as defined above, a client with an address) to the content provider. (That is, in the forward direction). In some embodiments, diverting the forward flow is an application-specific problem. Web traffic is relatively amenable to forward diversion. In non-Web
15 cases, especially for non-TCP protocols, forward diversion involves, in some embodiments, a special case analysis. In some embodiments, specific non-Web applications will be identified and optimized separately. In some embodiments, the applications of interest can be altered as to make forward flow diversion easier. The general model proposed involves a specialized device that proxies
20 for the application in question. In some embodiments, this proxy can be established on the fly. In other embodiments, this proxy will need manual configuration.

In this section, we describe a method for Web applications: this method involves altering the DNS name embedded in the URL the Surfer is accessing.
25 In some embodiments, this alteration can be made computationally or statically by the Content Provider, based on the logged in identity of the Surfer. The DNS name should be modified on the left (e.g., from www.foo.com to optimized.www.foo.com.)

30 In general, the Content Provider will normally own the authoritative DNS server for foo.com. In some embodiments, the Content Provider will insert a referral record for the subdomain "optimized" that has been added. The

referral will point to a specialized DNS server, which, in some embodiments, resides in the Bender, and, in other embodiments, resides in a CoLo. In some embodiments, the flexibility to return the normal answer for www.foo.com can be exercised; this could be appropriate in case a usage-based SLA has been
5 signed with the CP, and they have reached their limit on how much traffic they can optimize.

In other embodiments, the forward flow can be bended, so as to simplify the selection of return traffic to accept into the optimized routing fabric. In such embodiments, the Bender may no longer need to advertise a /32 to the Content
10 Provider for flows of interest; instead, all to-be-optimized flows can be forced back to the Bender, and no extra flows will be delivered (a potential occurrence with /32 advertisement). For Web traffic at least, the volume of forward flows is substantially less than the returns (most anecdotal measurements imply a 1:6 to 1:8 ratio in volume); hence, in such embodiments, all forward top-be-optimized
15 flows could be accepted without undue threat to the capacity of the core infrastructure.

In some embodiments, the commitment to send the forward traffic across the optimized routing fabric does not carry a reverse requirement; in some
20 implementations, so long as the Bender or CoLo rack can immediately arrange a packet which the Internet can return to the Surfer, and the Surfer can accept as valid, then the freedom can be kept to only take those return flows where a route with better performance exists.

In the implementation described above, the Web-based approach can
25 involve advertising a DNS response to qualified surfers (potentially selected by the Content Provider); this response is a special address, so that the Surfer sends its forward packets in via one of the PoPs. In some embodiments, this address can route all the way through to the Bender location across the optimized routing fabric (using the optimized PoP to PoP routes described earlier).

At the Bender, different implementations exist, depending, for example,
30 on the constraints of the Content Provider. In some embodiments, a combination of destination NAT (as commonly implemented in load balancing products) to forward the packet to the CP's Server, and source NAT (as

commonly implemented in firewall products, as well as IOS) can be used to ensure the response packets will return to the Bender. In some implementations, this will anonymize the traffic (the Content Provider no longer sees the original source address)

5 In other embodiments, the Content Provider will add a secondary address to their Server (or Server farm, represented by a Load Balancer's VIP address), and a method to route packets from the Bender to that isolated special address. In this case, no NAT is required. In some embodiments, a return to the original Bender formulation with /32 advertisements will be needed. The
10 advantage of such an implementation includes the enhanced ability of the Bender to identify properly qualified traffic; anything incorrectly passing through the Bender will not have the right source address from the Server.

In some embodiments, where there is a concern about unauthorized access to the Optimized Routing Fabric, this URL/DNS-based technique
15 techniques can be added to protect from such occurrences. In some embodiments, such techniques involve the Content Provider to notify that a Surfer is qualified (using the forward bending technique described above). Those skilled in the art will identify other such techniques. In some embodiments, when this technique is used in addition to the independent return flow mechanism described at length here, the Content Provider could explicitly
20 identify the qualified addresses and, in some embodiments, their "Time To Live" (TTL) value. In some embodiments, this technique can be combined with DNS redirection, in such a way that flows can be checked as they arrives at the Bender. In some embodiments, in case bottleneck avoidance needs to be enhanced, a flow can be permitted by default, but asynchronously looked up in
25 the currently qualified list. In such an implementation, a flow that proves to be unqualified can be simply redirect. In some embodiments, in case the abuse can be proven to be intentional, a TCP RST can be forced into the stream to punish the Surfer. In some embodiments, given the asymmetry of Web traffic,
30 disqualified forward traffic can be simply tolerated, and a negative rule can be added to the Bender's decision process on which flow to take over the optimized routing fabric. (In such an implementation, the unauthorized user can get no benefit on return traffic.) In other embodiments, the forward traffic can

be pushed off the optimized routing fabric. The technique could require, in some implementations, the intervention from the Content Provider, who could flag the address to those surfers, and potentially ask them to use HTML redirection to guide the Surfer back to the usual non-optimized Internet.

5 Relating to an embodiment of Figure 2, suppose we have a flow at point 1 trying to reach internetnetwork2, and the paths have the metrics associated with them as annotated, where lower is better and metrics can be added together. Some possibilities for determining the exit point and path to the exit point include (but are not limited to): choose the exit point first and then
10 determine path to the exit point. In this case, choose exit point 2 as the best exit point because its path to internetnetwork2 has metric 100 and exit point 3 has metric 200. Given that we have chosen exit point 2, we can choose either the best direct path from point 1 to exit point 2 (500) or the best indirect path (in this case, via exit point 3, 100+200). This results in a total path metric
15 of $100+200+100=400$. Consider both the exit point's paths and the paths to the exit point in combination. In this case, choose exit point 3, because even though the path to internetnetwork2 from exit point 3 is 200, the path to exit point 3 from point 1 is only 100, resulting in a total path metric of $100+200=300$.

20 Relating to one embodiment, Figure 3a shows an internetnetwork1 with a network device can decide to redirect flows directly bound for internetnetwork 4 to another internetnetwork with a network device, like internetnetwork 2, or to an internetnetwork without a network device, like internetnetwork 3.

25 Relating to one embodiment, Figure 3b shows an internetnetwork1 can have both flow endpoints inside of it, and can have multiple network devices deployed to obtain measurements, calculate metrics, select paths, and alter routing in the internetnetwork.

Measurement Packets

30 A measurement packet is a packet sent by a sender over an internetnetwork that includes information necessary for the receiver of the packet to compute measurements of the performance characteristics of the path the packet has traversed over that internetnetwork. The information includes information for a

receiver of the measurement packet to compute measurements of performance characteristics of at least a portion of the path of the measurement packet; and data including one or more of measurement statistics, a generic communication channel, network information, and control data directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

In some embodiments of the invention, the information included in the measurement packet to compute measurements includes at least one of a timestamp of a sending time of the packet and a number to identify the packet by itself and/ to identify the relative position of the measurement packet in a sequence of measurement packets,

In some embodiments of the invention, the measurement packet is implemented using the following data structure:

```

struct MeasurementHeader {

    /**
     * A generation number. This value represents when the
     * sender began sending. This value is a standard Unix
     * timestamp that seconds since Jan 1, 1970 UTC.
     */
    uint32_t mGeneration;

    /**
     * A sequence number for the packet. This increments each
     * time a packet is sent and rolls over when 16 bits is
     * exceeded.
     */
    uint16_t mSequence;

    /**
     * The IP address the packet is sent to.
     */
    uint32_t mDstAddr;

    /**
     * The send timestamp for this packet.

```

```

        **/
        uint64_t mSendTime;
    };

```

5 The mGeneration field is used to detect when a sending process has started a new session. This field is used by the receiver to determine that a discontinuity in the stream's sequence numbers is the result of a sender restart, rather than due to large network latencies, duplicate packets or dropped packets.

10 The sequence number mSequence field is incremented by one each time a packet is sent. This approach allows the receiver to deduce lost and duplicate packets by identifying missing and duplicate sequence numbers.

15 The mSendTime field contains the time at which the packet was sent, represented as microseconds since January 1, 1970 UTC. This field is compared to the time the packet arrived at the receiver to determine the delay between the sender and the receiver.

20 In some embodiments of the invention, a plurality of one or more packets are sent over a path continuously. In some embodiments of the invention, the continuous stream of packet is denoted as a measurement stream. Each measurement stream is uniquely identified by the source and destination IP addresses. The sender maintains one socket descriptor for each source IP address it sends from and writes the destination IP address into the mDstAddr field. On the receiver side, the source IP address is returned by the recv() system call and the destination address is retrieved from the measurement packet.

25

Data Included in the Measurement Packets

30 In measurement packets that contain sufficient space, data will be included, including one or more of measurement statistics, a generic communication channel, network information, and control data directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

Some embodiments of the invention will add a single type of data to each packet. Some embodiments of the invention will use a complex data, including subpackets.

Some embodiments of the invention use subpackets that include a single byte subpacket type identifier, followed by a 2-byte length field (including the length of the type and length fields) and finally including the data that is to be sent. One embodiment will store all values in network byte order. Other byte orders will be apparent to those skilled in the art. The following data structure definition describes some embodiments.

```
class SubPacket {  
    /*  
     * The type identifier for this subpacket.  
     */  
    uint8_t mType;  
  
    /*  
     * The length of this subpacket, in network byte order.  
     */  
    uint16_t mLength;  
};
```

One embodiment of this invention will include data describing a momentary snapshot of the measurement statistics for a given path between a sender and a receiver.

In some embodiments of this invention, this data will include one or more of the following information: the source and destination IP addresses that define the path, a measurement packet size for which the statistics have been calculated as well as computed measurement statistics that are at least partly responsive to delay; computed measurement statistics that are at least partly responsive to jitter and computed measurement statistics that are at least partly responsive to packet loss.

In one embodiment of this invention, these statistics will be in units of microseconds expressed as 64-bit floating-point quantities and transmitted in a standard network byte order.

In one embodiment of this invention, the following data structure will store the computed statistics:

```
class TunnelStatsSubPacket : public SubPacket {
5      /**
        * The time that this statistic snapshot was taken (in
        * microseconds since 1970).
        */
        uint64_t mTimestamp;
10
        /**
        * The source IP address of the tunnel these statistics
        apply
        * to.
        */
15      uint32_t mSrcAddr;

        /**
        * The destination IP address of the tunnel these
20      statistics
        * apply to.
        */
        uint32_t mDstAddr;

25      /**
        * The size of measurement packet that these statistics
        apply
        * to. A size of 0 indicates that these statistics
        apply to
30      * all packet sizes.
        */
        uint16_t mPktSize;

        /**
35      * The average delay in microseconds.
        */
        double mDelay;

        /**
```

```

        * The average jitter in microseconds.
        **/
double mJitter;

5      /**
        * The percentage of packets that have been lost, in the
range
        * 0 to 1.
        **/
10     double mLoss;
};

```

Some embodiments of this invention include the time at which the statistics were computed such that those statistics are sent over multiple paths for improved reliability and to take advantage of one path having less delay than another. One embodiment at the receiving end is able to compare the computation times of received statistics to place them in their original temporal order, regardless of their relative arrival times over the paths.

Some embodiments of this invention will send computed statistics specific to the paths that are part of the plurality of one or more paths that are between the specific sender and receiver. Other embodiments will send additional computed statistics for paths that are not one of the plurality of one or more paths that are between the specific sender and receiver.

Some embodiments of this invention will include network information concerning network topology including but not limited to information retrieved from routers such as in-bound or out-bound link utilization, inbound or out-bound link bandwidth and/or CPU utilization. Other network information determined from routers and other network devices will be apparent to someone skilled in the art.

Some embodiments of this invention will also include control data directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

In some embodiments of the invention, the control data will instruct a receiver to alter its configuration, including but not limited to zero or more of the following examples: instructing a receiver to initiate sending a plurality of

one or more measurement packets, change one or more of the measurement packet sizes, inter-measurement packet transmission times and mix of packet sizes, and stop sending one or more of the plurality of measurement packets.

5 In some embodiments of the invention, this control information will include notification of measurement devices that have joined or left the network.

In many embodiments of the invention, the measurement packets will be encrypted by the sender and decrypted by the receiver. Some of these embodiments will use IPSec.

10 In some embodiments of the invention, the encryption and decryption will be done by an external device using IPSec.

Other encryption and decryption options will be apparent to one skilled in the art.

15 In some embodiments of the invention, the measurement packets will be digitally signed.

In some embodiments of the invention, a generic communication channel will be used by a sender and a receiver to communicate data between them.

20 **Performance Characteristics of a Path**

Measurements are used to compute performance characteristics of the paths traversed by the measurement packets. The measurements can either be computed from the measurement packets themselves, or extracted from the arbitrary data carried by the measurement packets. The measurements of
25 performance characteristics include at least one or more of one-way measurements and round-trip measurements. The performance characteristics include at least one or more reachability, delay, jitter, loss, available bandwidth, and total bandwidth. Other performance characteristics will be apparent to those skilled in the art.

30 In some embodiments of the invention, delay measurements are computed as the interval of time from the moment the measurement packet is sent by the sender to the moment of time the measurement packet is received by the receiver. The sending time is carried by the packet, and it is measured by the

clock the sender refers to. The receiving time is measured by a clock that the receiver refers to, which may or may not be synchronized with the sender's clock.

5 In some embodiments of the invention, the clock of the sender and the clock of the receiver are synchronized. A plurality of one or more precise clock inputs such as GPS, NTP, IRIG and NIST will be used. Some embodiments of this invention will use the same clock as an input to more than one of the plurality of one or more senders and receivers. In some embodiments of the invention, the clock of the sender and the clock of the receiver are the same.

10 In some embodiments of the invention, the clock of the sender and the clock of the receiver are not synchronized, and mechanisms based on the measurement data are used to correct the clock skew and clock drift, the mechanisms including using minimum delay across multiple measurement samples, and using a mechanism to track the minimum delay over time.

15 Some embodiments of the invention will use the minimum round-trip delay between the devices to place a lower bound on clock skew.

Some embodiments of the invention will use the lower bound of multiple paths between the sender and receiver to further reduce the lower bound.

20 Some embodiments of the invention will correct for clock drift by tracking the relative clock skew between the sender and receiver over time and adjusting for the slope of the drift.

In some embodiments of the invention, jitter measurements, also known as inter-measurement packet delay variations, are computed as the difference in delay on consecutive, successfully received packets.

25 In some embodiments of the invention, jitter can also be computed as the difference between the instantaneous delay of a packet, and the average delay of all the measurement packets previously received.

30 In some embodiments of the invention, loss measurements are computed by assigning a timeout value to each measurement packet that indicates the instant of time after which the measurement packet will be declared lost, if the packet has not arrived by that time. In some embodiments of the invention, the timeout value of a measurement packet can be computed with the transmission

time of a previously received packet, an estimation of the inter-transmission time between measurement packet, and an estimation of the transmission delay of the measurement packet. In some embodiments of the invention, the inter-transmission time can be estimated if the receiver knows about the scheduling pattern of transmission of measurement packets. In some embodiments of the invention, the transmission delay of packet can be estimated based on delay and jitter performance characteristics.

Performance characteristics of a path could be the measurement themselves, or statistics on those measurements. In the statistics case, a dynamic algorithm is used to update the statistics associated with a path with every new measurement obtained with the arrival of every new packet over the path.

In some embodiments of the invention, the algorithm computes statistics over the performance characteristics of the path.

In some embodiments of the invention, the statistics include averages, deviations, and variances. Other statistics will be apparent to those skilled in the art. In some embodiments of the invention, averages can be computed using a plurality of one or more techniques including a moving average, an average based on the Robbins-Moro estimator, a window-based average or a bucket-based average. Other techniques to compute averages will be apparent to those skilled in the art.

In some embodiments of the invention, the moving average is an exponentially moving average computed using a Robbins-Moro estimator. The Robbins-Moro stochastic approximation estimator finds a solution of the equation:

$$E[f(t) - x] = 0$$

where E is the expectation, f(t) a function and x the estimator. The general form of the solution is:

$$x(t) = x(t-1) + \alpha * [f(t) - x(t-1)] = (1 - \alpha) * x(t-1) + \alpha * f(t)$$

or, with $\alpha = (1 - \mu)$,

$$x = \mu * x + (1 - \mu) * f$$

μ is the weight of the estimator, and determines the amount contributed to the average by the function.. In some embodiments of the invention, μ is constant. In some embodiments of the invention, μ is a dynamic value, whose value depends on the last value of the function f according to the formula:

$$\mu = e^{(-f/K)}$$

where K is a constant that also determines the importance of the last value of f with respect to the current value of the estimator x .

In some embodiments of the invention, average delay can be computed using an exponentially moving average as follows,

$$d = \mu * d + (1 - \mu) * m$$

where d is the exponentially moving average of delay, m is the last delay sample, and μ is the weight of the moving average.

In some embodiments of the invention, average jitter can be computed using an exponentially moving average as follows,

$$v = \mu * v + (1 - \mu) * |d - m|$$

where v is the exponentially moving average of jitter, $|d - m|$ is the last sample of jitter, and μ is the weight of the average.

In some embodiments of the invention, average jitter can be computed using an exponentially moving average as follows,

$$v = \mu * v + (1 - \mu) * |m - m'|$$

Where v is the exponentially moving average of jitter, $|m - m'|$ is the last sample of jitter, m is the last delay sample, m' is the previous delay sample, and μ is the weight of the average.

In some embodiments of the invention, delay and jitter averages can be combined into a single value as follows:

$$l = d + M * v$$

- 5 Where d is the average delay, v is the average jitter and M is a constant.

In some embodiments of the invention, average loss can be computed using an exponentially moving average as follows,

$$p\text{-hat} = \mu * p\text{-hat} + (1 - \mu) * p$$

- 10 where $p\text{-hat}$ is the moving average of the loss, $p = \{0 \text{ if packet is received, } 1 \text{ if the packet is declared lost}\}$, and μ is the weight of the exponentially moving average.

- In some embodiments of the invention, μ is determined based on the notion of forgiveness against a single packet loss. The forgiveness period is the interval of time between the time the packet loss occurs and the time the average loss is forgiven. The forgiveness period can be either defined in units of time, or in number of packets if the rate of the monitoring flow is known. That is, the forgiveness period will end after n consecutive packets have been received after the loss, when these packets have been transmitted at a certain rate.
- 15
20

- The value of the exponentially moving average after receiving the n packets is needed before μ can be determined, and this value is known as the forgiveness threshold. In some embodiments of the invention, the forgiveness threshold is chosen arbitrarily. In some embodiments of the invention, the forgiveness threshold takes the value:
- 25

$$\frac{1}{2} (1 - \mu)$$

This value is half of the value of the estimator after the single loss occurs, and thus we call it the *half-life threshold*. Similarly, we also call the forgiveness period under this threshold the *half-life period*. The advantage of using a

forgiveness threshold greater than zero is that issues related to host-dependent floating-point representations reaching that value are avoided.

In some embodiments of the invention, μ is computed by comparing the value of the estimator after n consecutive packet arrivals since the loss with the *half-life threshold*:

$$p\text{-hat} = (1 - \mu) * \mu^n < \frac{1}{2} (1 - \mu)$$

Given that n is known because is determined by the value of the *half-life period* and the transmission rate, μ is computed as:

$$\mu = \exp((\ln \frac{1}{2}) / n)$$

In some embodiments of the invention, two thresholds are defined, an upper threshold and a lower threshold. When the value of $p\text{-hat}$ exceeds the upper threshold, the loss is not forgiven until enough measurement packets are received consecutively so that the value of $p\text{-hat}$ gets below the lower threshold.

Other mechanisms to compute μ will be apparent to for those skilled in the art.

Path Description

In some embodiments of the invention, the path traversed by the measurement packets from the sender to the receiver is such that the path is at least partly implemented with at least one of a GRE tunnel, an IPSEC tunnel and IPonIP tunnel. Other path implementations using tunnel will be apparent for those skilled in the art.

In some embodiments of the invention, the path traversed by the measurement packets from the sender to the receiver is implemented with a virtual circuit, including a frame relay PVC, an ATM PVC or MPLS. Other path implementations using virtual circuits will be apparent for those skilled in the art.

Other path implementations will be apparent to those skilled in the art.

Internetwork Description

In some embodiments of the invention, the internetwork is implemented by a plurality of one or more subnetworks, including a plurality of one or more

VPNs, a plurality of one or more BGP autonomous systems, a plurality of one or more local area networks, a plurality of one or metropolitan area networks, and a plurality of one or morewide area networks.

In some embodiments of the invention, the internetwork is implemented by an overlay network.

Other internetwork implementations will be apparent to those skilled in the art.

Packet Sizes and Transmission Times

In some embodiments of the invention, the measurement packets are of varying sizes, including 64, 256, 512, 1024, 1500 bytes.

In some embodiments of the invention, the size of the measurement packets is specified with an external API.

In some embodiments of the invention, the measurement packets are of a fixed size.

In some embodiments of the invention, the measurement packet sizes and times between measurement packets simulate the traffic pattern of a plurality of one or more applications

In some embodiments of the invention, traffic patterns correspond to voice applications, where the packets re of small size, e.g., 30 bytes, and the inter-transmission time between consecutive packets is constant, e.g., 10 ms. These examples do not limit the possible size values and inter-transmission time values.

In some embodiments of the invention, traffic patterns correspond to video applications, where the packets size is the largest permitted to be transmitted by an internetwork without being fragmented, and the inter-transmission time between consecutive packets varies depending on the spatial and temporal complexity of the video content being transmitted, the compression scheme, the encoding control scheme.

In some embodiments of the invention, traffic patterns correspond to the plurality of applications observed in an internetwork, including at least one or more of HTTP transactions, FTP downloads, IRC communications, NNTP exchanges, streaming video sessions, VoIP sessions, videoconferencing sessions

and e-commerce transactions. Other types of applications will be apparent to those skilled in the art.

In some embodiments of the invention, the inter-measurement packet transmission times are of varying length.

5 In some embodiments of the invention, the inter-measurement packet transmission times are of fixed length.

In some embodiments of the invention, the inter-measurement packet transmission times specified with an external API.

10 In some embodiments of the invention, the length of the inter-measurement packet transmission times is randomized according to a distribution. In some embodiments of the invention, this distribution is based at least in part on a uniform distribution. In some embodiments of the invention, this distribution is based at least in part on an exponential distribution. In some
15 embodiments of the invention, this distribution is based at least in part on a geometric distribution. Other distributions will be apparent to those skilled in the art.

In some embodiments of the invention, the length of the inter-measurement packet transmission times is provided by a table.

20 In some embodiments of the invention, the length of the inter-measurement packet transmission times is controlled by a scheduler. In some embodiments of the invention, the scheduler uses a priority queue, keyed on desired send time.

25 Other mechanisms to specify the inter-measurement packet transmission time will be apparent to those skilled in the art.

Other packet sizes and transmission times will be apparent to those skilled in the art.

Path Selection

30 It is possible that multiple alternative paths between a sender and a receiver are available through an internetwork at any given moment. Performance characteristics of each of these paths can be used to select a subset of the paths.

In some embodiments of the invention, the subset of the plurality of paths is selected based at least in part on at least one of: one or more of the measurement statistics from the measurement packet and one or more of the computed statistics.

5 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the position of paths in a ranking. In some embodiments of the invention, the ranking is at least partly based on one or more of the measurement statistics included as data in the measurement packet. In some embodiments of the invention the ranking is at least partly
10 based on the computed statistics of the path. In some embodiments of the invention the ranking is implemented by using a comparison function to compare the paths, and by ordering the paths in a decreasing order. In some embodiments of the invention the ranking is implemented by using a comparison function to compare the paths, and by ordering the paths in an
15 increasing order. Other ranking techniques will be apparent to those skilled in the art.

In some embodiments of the invention, the ranking is based on a single score associated to each path. In some embodiments of the invention, this score is denoted *Magic Score* (MS), and it is computed as follows:

20

$$MS = ML * MF$$

$$ML = d + M * \nu$$

$$MF = \text{delta} * p\text{-hat} + 1$$

where ML is the *Magic Latency*, a component of the MS obtained using delay and jitter respectively calculated with statistics; and MF is the *Magic scaling Factor* that multiplies the value of ML, and is computed based on loss statistics. M is a constant that takes several values, including 4, for example. MS can be seen as a scaled-up version of ML, and the scaling factor MF is a function of p-hat and delta, a constant. As p-hat not only reflects loss but also detects large
25 delay spikes before they happen, p-hat can be seen as an indicator of the departure of the path from a "normal mode" operation, and thus the scaling
30

factor is only applied when there are loss or spikes. The goal of MF is to differentiate between paths that have very similar delay characteristics, but with one having losses and the other not having them.

5 In some embodiments of the invention, ML is used as a delay indicator, given that jitter is accounted as an increase in delay. In contrast, MS, although a scaled version of ML, cannot be used to indicate delay, except when $MF = 1$ ($\hat{p} = 0$), which leads to $MS = ML$. That means the value of MS is useful not by itself but to compare it with the MSs of other tunnels.

10 In some embodiments of the invention, loss statistics can be used as a discriminator instead of a scaling factor. That is, \hat{p} can eliminate paths experimenting loss. Then, the remaining paths can be selected using $MS = ML$.

In some embodiments of the invention, the selection of a subset of paths is based on applying at least one or more thresholds to at least one of more of the statistics.

15 In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.

20 In some embodiments of the invention, the number of paths in the subset of paths is fixed. In some embodiments of the invention, this fixed number of paths N out of M paths is determined such that the probability of having loss in $(M - N)$ paths simultaneously is less than a certain threshold. In some embodiments of the invention, this probability is a binomial, with the assumption that all paths have the same probability of loss.

25 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on a probability associated with each path. In some embodiments of the invention, the probability of each path is at least partly based on one or more of the measurement statistics included as data in the measurement packet.

30 In some embodiments of the invention, the probabilities of each path are equal.

In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the cost of the path.

5 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the amount of bandwidth consumed over a period of time.

Other possibilities to compute path probabilities will be apparent to those skilled in the art.

Other mechanisms to select a subset of the paths will be apparent to those skilled in the art.

CLAIMS

1. A method of optimizing traffic in an internetwork, the method comprising:

5 selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each
10 of the subset of flows, the measuring performance characteristics including one or more of:

launching a first plurality of one or more packets, and measuring the first plurality of one or more packets;

measuring a second plurality of one or more packets in the internetwork,
15 wherein the second plurality of one or more packets were already launched;

monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more
20 flows, and measuring the second plurality of one or more flows; and

encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a
25 second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the
30 at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and

in response to calculating the at least one performance metric,
rearchitecting the internetwork to optimize one or more of the at least one
performance metric, rearchitecting the internetwork including at least one of:

altering a plurality of one or more routing tables in the internetwork,

wherein the plurality of one or more routing tables include at least one of:

network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2
forwarding tables, and MPLS forwarding tables;

redirecting the subset of flows to a second internetwork coupled
to the internetwork at one or more Points of Presence; and

affecting forwarding decisions of the subset of flows, by
imposing one or more of: NAT, GRE, and tunneling techniques other
than GRE.

2. The method of claim 1, wherein the plurality of one or more routing
tables includes network-layer routing tables.

3. The method of claim 1, wherein the plurality of one or more routing
tables includes layer 3 routing tables.

4. The method of claim 1, wherein the plurality of one or more routing
tables includes IP routing tables.

5. The method of claim 1, wherein the plurality of one or more routing
tables includes layer 2 routing tables.

6. The method of claim 1, wherein the plurality of one or more routing
tables includes MPLS forwarding tables.

7. The method of claim 1, wherein the performance characteristics includes
the plurality of one or more round trip measurements.

8. The method of claim 1, wherein the performance characteristics includes
the plurality of one or more one-way measurements

9. The method of claim 1, wherein the one or more of the at least one performance metric to be optimized quantifies path performance of the subset of flows.

10. The method of claim 9, wherein the one or more of the at least one performance metric characterizes a quality of a network application.

11. The method of claim 10, wherein the application includes http.

12. The method of claim 10, wherein the application includes ftp.

13. The method of claim 10, wherein the application includes telnet.

14. The method of claim 10, wherein the application includes voice.

15. The method of claim 10, wherein the application include video.

16. The method of claim 1, wherein one or more of the at least one performance metric quantifies one or more of: overall usage of the at least the portion of the internetwork, absolute individual network link usage in the at least the portion of the internetwork, relative individual network link usage in the at least the portion of the internetwork, and link usage cost in the at least the portion of the internetwork.

17. The method of claim 1, wherein the subset of flows include at least a first sub-plurality of one or more flows, and a second sub-plurality of one or more flows, and the first sub-plurality of one or more flows uses at least a first measured performance characteristic, and the second sub-plurality of one or more flows uses at least a second measured performance characteristic.

18. The method of claim 1, wherein the one or more of the at least one performance metric quantifies geographic distance covered by the subset of flows in the internetwork.

19. The method of claim 1, wherein the selection of the subset of flows is user-based in the at least the portion of the internetwork.

20. The method of claim 1, wherein the internetwork comprises an autonomous sub-system of a larger network.

21. The method of claim 20, wherein the larger network is the Internet.

22. The method of claim 20, wherein the internetwork is a BGP autonomous system.

23. The method of claim 1, wherein the internetwork comprises an autonomous system of a larger network.

24. The method of claim 23, wherein the larger network is the Internet.

25. The method of claim 23, wherein the internetwork is a BGP autonomous system.

26. The method of claim 1, wherein the internetwork comprises an overlay network.

27. The method of claim 1, wherein the internetwork comprises a plurality of one or more networks, the plurality of one or more networks coupled together, wherein the plurality of one or more networks include one or more of local-area networks, metropolitan-area networks, and wide-area networks.

28. The method of claim 1, wherein the measuring performance characteristics includes launching a first plurality of one or more packets, and measuring the first plurality of one or more packets.

29. The method of claim 28, wherein the measuring performance characteristics includes launching a first plurality of one or more round-trip packets, and measuring the first plurality of one or more round-trip packets, such that the plurality of one or more round trip measurements include round trip measurements for the launched packets.

30. The method of claim 29, wherein the first plurality of one or more round-trip packets include traceroute ICMP packets.

31. The method of claim 29, wherein the first plurality of one or more round-trip packets include ping ICMP packets.

32. The method of claim 29, wherein the first plurality of one or more round-trip packets include telnet packets.

33. The method of claim 29, wherein the first plurality of one or more round-trip packets include TCP packets from an empty TCP transaction.

34. The method of claim 29, wherein the first plurality of one or more round-trip packets include http packets.

35. The method of claim 28, wherein the measuring performance characteristics includes launching a first plurality of one or more one-way packets, and measuring the first plurality of one or more one-way packets, such that the plurality of one or more one-way measurements include one-way measurements for the launched packets.

36. The method of 1, wherein the measuring performance characteristics includes measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched

37. The method of claim 1, wherein round-trip measurements include one or more of round-trip delay, round-trip jitter, round-trip loss, round-trip available bandwidth, and round-trip total bandwidth.

5

38. The method of claim 37, wherein round-trip measurements include round-trip delay.

10

39. The method of claim 37, wherein round-trip measurements include round-trip jitter.

40. The method of claim 37, wherein round-trip measurements include round-trip loss.

15

41. The method of claim 37, wherein round-trip measurements include round-trip available bandwidth.

42. The method of claim 37, wherein round-trip measurements include round-trip total bandwidth.

20

43. The method of claim 1, wherein one-way measurements include one or more of one-way delay, one-way jitter, one-way loss, one-way available bandwidth, and one-way total bandwidth

25

44. The method of claim 43, wherein one-way measurements include one-way delay.

45. The method of claim 43, wherein one-way measurements include one-way jitter.

30

46. The method of claim 43, wherein one-way measurements include one-way loss.

47. The method of claim 43, wherein one-way measurements include one-way available bandwidth.

5 48. The method of claim 43, wherein one-way measurements include one-way total bandwidth.

10 49. The method of claim 1, wherein altering a plurality of one or more routing tables in the internetwork, wherein the plurality of one or more routing tables include at least one of: network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2 forwarding tables, and MPLS forwarding tables.

15 50. The method of claim 49, wherein the altering of the plurality of one or more routing tables is applied automatically.

20 51. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers, wherein the configuring the plurality of one or more routers statically alter the routing of flows.

52. The method of claim 51, wherein the configuring a plurality of one or more routers includes route maps.

25 53. The method of claim 51, wherein the configuring a plurality of one or more routers includes static route statements.

30 54. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers, wherein the configuring the plurality of one or more routers adjust the processing of dynamic routing updates.

55. The method of claim 49, wherein the altering of the plurality of one or more routing tables includes configuring a plurality of one or more routers,

wherein the configuring the plurality of one or more routers adjust the processing of dynamic routing updates.

56. The method of claim 49, wherein the altering of the plurality of one or more routing tables is performed at least partly by a plurality of one or more dynamic routing protocols.

57. The method of claim 56, wherein the plurality of one or more dynamic routing protocols includes BGP.

58. The method of claim 56, wherein the plurality of one or more dynamic routing protocols includes a plurality of one or more IGP routing protocols.

59. The method of claim 58, wherein the plurality of one or more IGP routing protocols includes OSPF.

60. The method of claim 49, wherein the altering of the plurality of one or more routing tables are applied manually by a user.

61. The method of claim 1, wherein the rearchitecting of the internetwork includes redirecting the subset of flows to a second internetwork coupled to the internetwork at one or more Points of Presence.

62. The method of claim 61, wherein the redirecting of the subset of flows is at least partly across an exit point traversed by at least a portion of the subset of flows from the internetwork, wherein the at least one performance metric is optimized.

63. The method of claim 1, wherein the optimizing of the performance metric includes:

calculating at least one performance metric for a plurality of one of more paths in the at least the portion of the internetwork; and

at least partly responsive to the calculating of the at least one performance metric, selecting a path in the at least the portion of the internetwork from the plurality of one or more paths.

5 64. The method of claim 63, wherein the selected path is a direct path.

65. The method of claim 63, wherein the selected path is an indirect path.

10 66. The method of claim 1, wherein the optimizing of the performance metric includes:

selecting a plurality of one or more exit points; and

selecting a plurality of one or more paths to reach the plurality of one or more exit points in the at least the portion of the internetwork.

15 67. The method of claim 66, wherein the steps of selecting the plurality of one or more exit points and selecting the plurality of one or more paths are performed separately

20 68. The method of claim 66, wherein the steps of selecting the plurality of one or more exit points and selecting the plurality of one or more paths are performed together

25 69. The method of claim 1, wherein the measuring performance characteristics includes monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets.

30 70. The method of claim 1, wherein the measuring performance characteristics includes serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows.

71. The method of claim 1, wherein the measuring performance characteristics includes encoding material within a third plurality of one or more

flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows.

72. The method of claim 1, wherein the measuring performance characteristics is done at least partly using flow information export.

73. The method of claim 1, wherein the measuring performance characteristics is done using RMON II.

74. The method of claim 1, wherein the measuring performance characteristics is provided by a source external to the subset of flows.

75. A method of optimizing traffic in an internetwork, the method comprising:

selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics includes one or more of:

measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;

monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and in response to calculating the at least one performance metric, affecting the routing of the subset of flows by altering a plurality of one or more DNS entries in the internetwork.

76. A network systems, comprising:

a plurality of one or more network devices configured such that, when the plurality of one more network devices is deployed in an internetwork, the plurality of one or more network devices performs:

selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics including one or more of:

launching a first plurality of one or more packets, and measuring the first plurality of one or more packets;

measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;

monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and

encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and

in response to calculating the at least one performance metric, rearchitecting the internetwork to optimize one or more of the at least one performance metric, rearchitecting the internetwork including at least one of:

altering a plurality of one or more routing tables in the internetwork, wherein the plurality of one or more routing tables include at least one of: network-layer routing tables, layer 3 routing tables, IP routing tables, layer 2 forwarding tables, and MPLS forwarding tables;

redirecting the subset of flows to a second internetwork coupled to the internetwork at one or more Points of Presence; and

affecting forwarding decisions of the subset of flows, by imposing one or more of: NAT, GRE, and tunneling techniques other than GRE.

77. The network system of claim 76, where the rearchitecting of the internetwork is performed within the device.

78. A network systems, comprising:

5 a plurality of one or more network devices configured such that, when the plurality of one more network devices is deployed in an internetwork, the plurality of one or more network devices performs:

selecting a subset of flows in the internetwork for monitoring, wherein the subset of flows includes one of one flow, some flows, and all flows;

10 measuring performance characteristics of the subset of flows in at least a portion of the internetwork, the performance characteristics including one or more of: a plurality of one or more round trip measurements for each of the subset of flows and a plurality of one or more one-way measurements for each of the subset of flows, the measuring performance characteristics including one
15 or more of:

measuring a second plurality of one or more packets in the internetwork, wherein the second plurality of one or more packets were already launched;

monitoring a first plurality of one or more flows in the internetwork, duplicating one or more packets from the plurality of one
20 or more flows, and measuring the one or more duplicated packets;

serving as a proxy hop for a second plurality of one or more flows, and measuring the second plurality of one or more flows; and

25 encoding material within a third plurality of one or more flows, causing a fourth plurality of one or more flows to be generated, wherein the third plurality of one or more flows traverses a first path of the internetwork, and the fourth plurality of one or more flows traverses a second path of the internetwork, wherein at least a portion of the first path of the internetwork and at least a portion of the second path of the
30 internetwork are equal, and measuring a subset of the fourth plurality of one or more flows;

calculating at least one performance metric for the subset of flows in the at least a portion of the internetwork, the at least one performance metric at least partly determined from the measured performance characteristics; and

in response to calculating the at least one performance metric, affecting the routing of the subset of flows by altering a plurality of one or more DNS entries in the internetwork

79. The network system of claim 78, where the affecting the routing of the subset of flows is performed within the device.

80. A method of populating a plurality of one or more groups with a plurality of one or more network addresses, the method comprising:

selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices;

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

81. The method of claim 80, wherein at least path performance is selected.

82. The method of claim 80, wherein at least network cost is selected.

83. The method of claim 80, wherein at least network usage is selected.

84. The method of claim 80, wherein at least geographical proximity is selected.

5 85. The method of claim 80, wherein at least topological proximity is selected.

86. The method of claim 80, wherein at least statistical similarity is selected.

10 87. The method of claim 80, wherein the at least the classification function includes a plurality of one or more criteria, the plurality of one or more criteria including at least one of topological proximity and administrative policy, wherein the topological proximity includes one or more of adjacency and dependency, and the administrative policy includes one or more of address aggregation and user-defined policy.

15 88. The method of claim 87, wherein the plurality of one or more criteria includes at least topological proximity, wherein the topological proximity includes one or more of adjacency and dependency.

20 89. The method of claim 87, wherein the plurality of one or more criteria includes at least administrative policy, wherein the administrative policy includes one or more of address aggregation and user-defined policy.

25 90. The method of claim 80, wherein, after the populating, the plurality of one or more groups includes a group, the group including network addresses corresponding to the plurality of one or more metrics having a value in a pre-defined range.

30 91. The method of claim 80, wherein, after the populating, the plurality of one or more groups includes clusters created by a clustering algorithm applied to at least one of the plurality of one or more metrics.

92. The method of claim 80, wherein the selecting, the creating, and the classifying are performed at a central server in an autonomous system.

93. The method of claim 80, wherein the selecting, the creating, and the classifying are performed at a plurality of one or more distributed servers distributed through a plurality of one or more autonomous systems.

94. The method of claim 80, wherein the selecting, creating, and populating are repeated, at least partly responsive to one or more of: a change in the at least one of the plurality of one or more metrics, an expiration of a time period, and an external input.

95. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the change in the at least one of the one or more metrics.

96. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the expiration of a time period.

97. The method of 94, wherein the selecting, creating, and populating are repeated, at least partly responsive to the external input.

98. The method of 94, wherein, prior to the change, a first plurality of one or more groups is populated, and after the change, a second plurality of one or more groups is populated, and the first plurality of one or more groups is at least partly different from the second plurality of one or more groups.

99. The method of claim 80, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from an observed routing table.

100. The method of claim 99, wherein the observed routing table is from an autonomous system, such that each of the plurality of network addresses is reachable by the autonomous system.

5 101. The method of claim 99, wherein the default population is a subset of the observed routing are performed on an existing routing table.

102. The method of 101, wherein after selecting, creating, and populating are performed on the existing routing table, any changes to the existing routing
10 table are limited to combining two or more entries into one group.

103. The method of 101, wherein after selecting, creating, and populating are performed on the existing routing table, changes to the existing routing table include dividing at least one entry into two or more groups.

15 104. The method of claim 80, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from observed network traffic.

20 105. The method of claim 80, wherein the selecting, creating, and populating are performed in a first autonomous system.

25 106. The method of claim 105, further comprising:
prior to populating the plurality of groups, creating a default population for at least one of the plurality of one or more groups, wherein the default population is derived at least partly from an observed routing table.

30 107. The method of 105, wherein the observed routing table is from the first autonomous system.

108. The method of 105, wherein the observed routing table is from a second autonomous system.

109. The method of claim 80, wherein at least one subset of the plurality of one or more groups includes contiguous addresses.

5 110. The method of claim 80, wherein at least one subset of the plurality of one or more groups includes at least two discontinuous addresses.

111. The method of claim 80, where the plurality of one or more network addresses includes one or more IP addresses.

10

112. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

15

responsive to 1) selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity, and 2) creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

20

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

25

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

30 113. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing

devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

114. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

responsive to selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity,

creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices;

populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

5 115. A network system, comprising:

a plurality of one or more network devices, the plurality of network devices including one or more of servers, forwarding devices, and routing devices, the plurality of network devices populating a plurality of one or more groups with a plurality of one or more network addresses by performing:

10 selecting a plurality of one or more metrics, the plurality of one or more metrics including one or more of path performance, network cost, network usage, geographical proximity, topological proximity, and statistical similarity;

responsive to creating the plurality of one or more groups in one or more network devices, network devices including one or more of servers, forwarding devices, and routing devices,

15 populating each of the plurality of one or more groups with a subset of the plurality of one or more network addresses based on at least a classification function, the classification function at least partly determined by at least one of the plurality of one or more metrics; and

20 including one or more network addresses from one or more groups of the plurality of one or more groups in a plurality of one or more routing tables distributed across the internetwork.

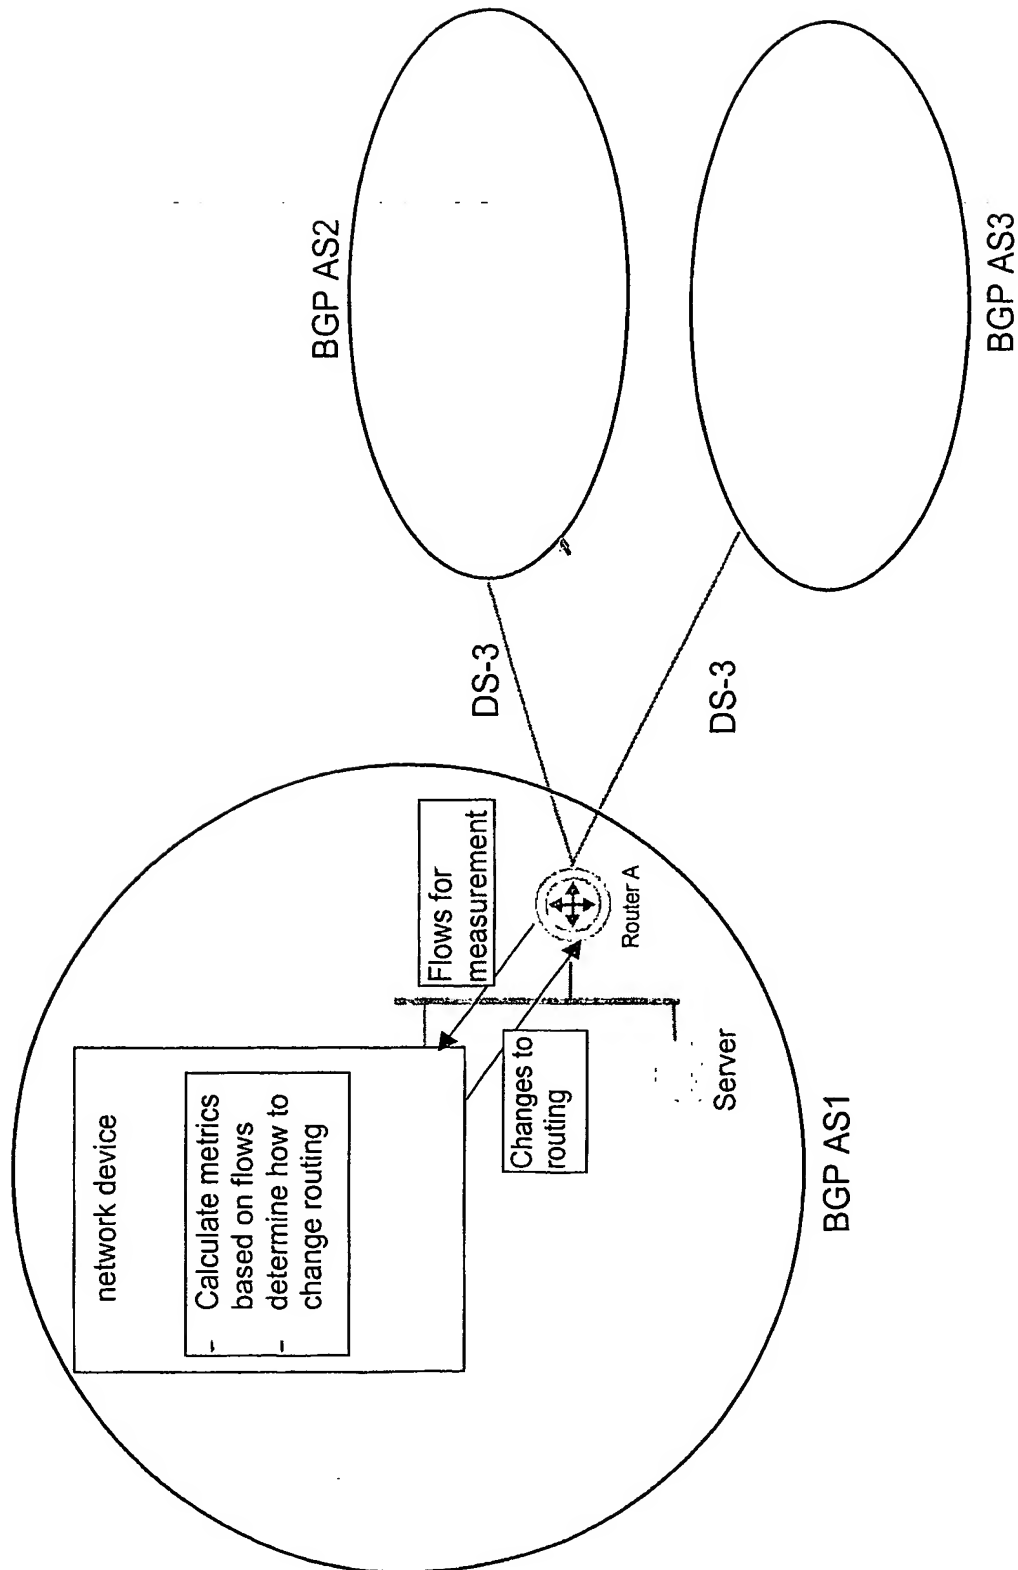


Figure 1




Figure 2

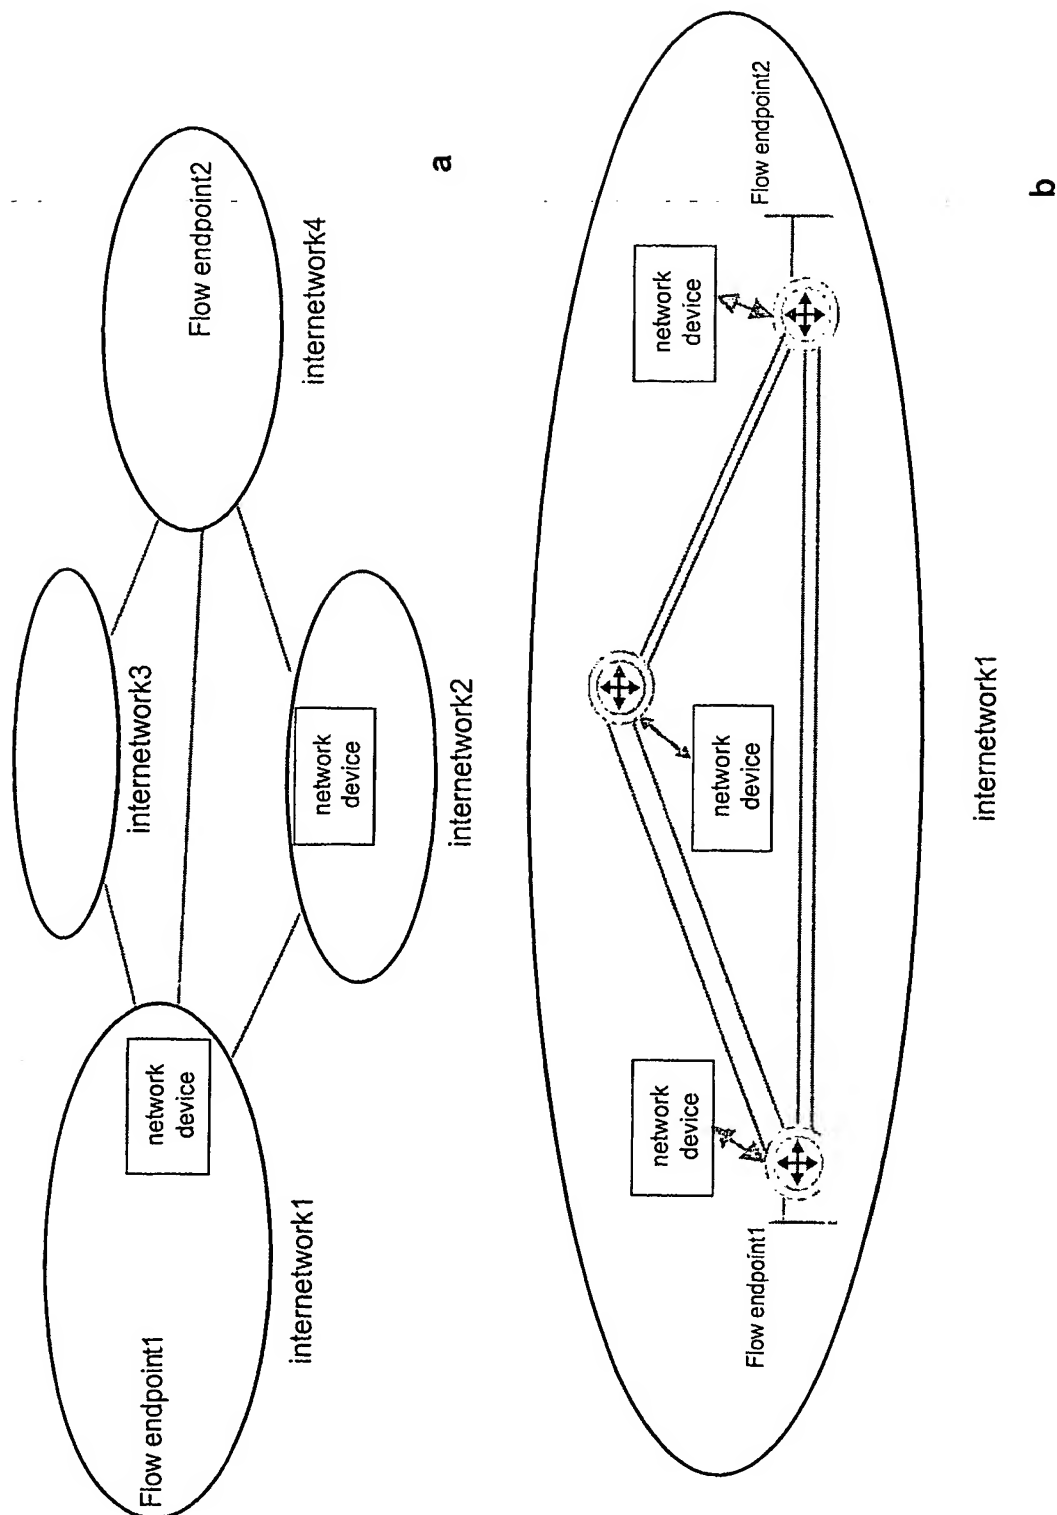


Figure 3

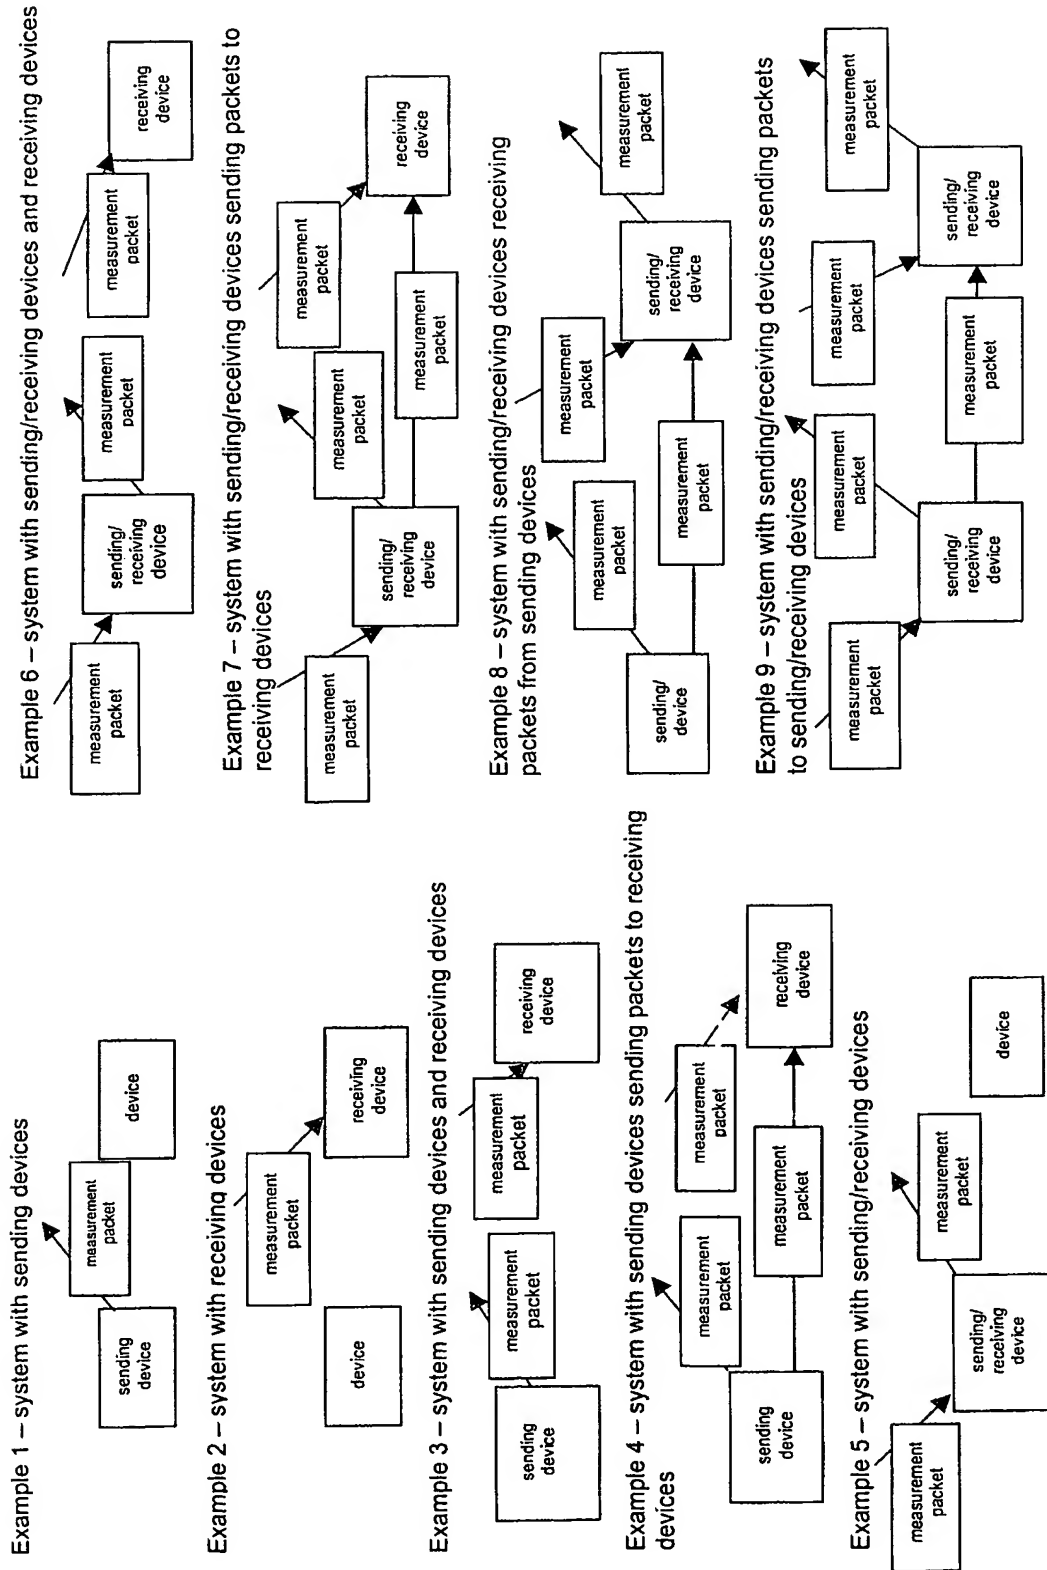


Figure 4

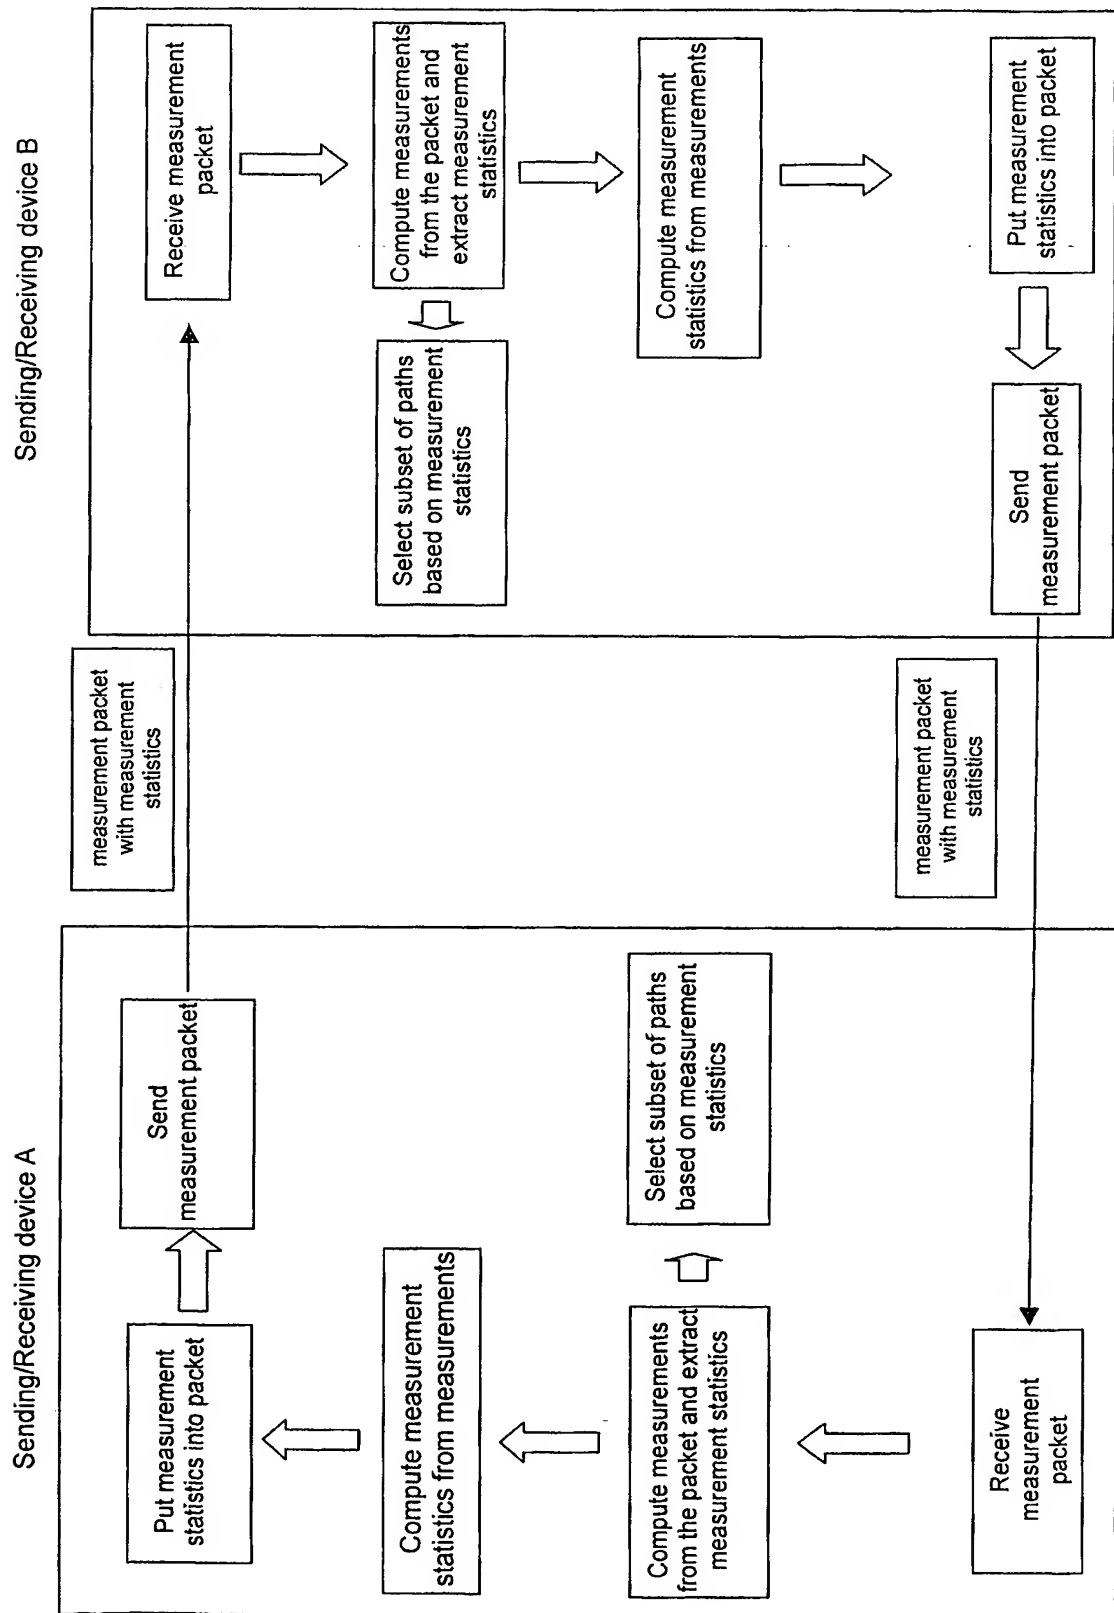


Figure 5

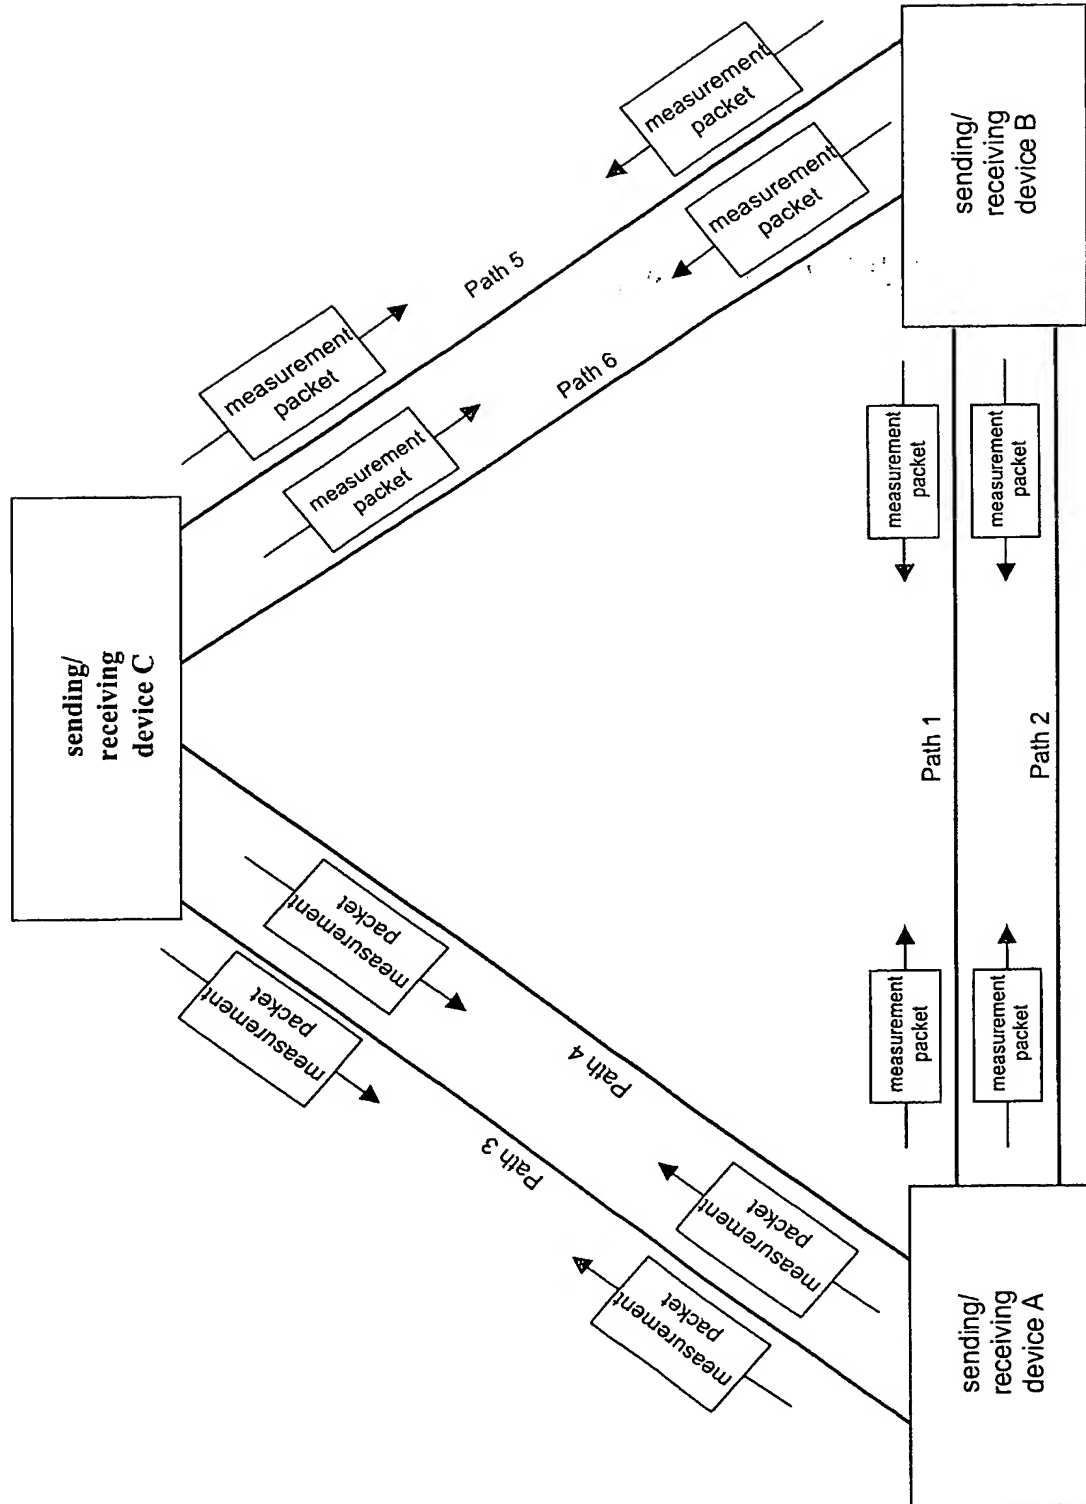


Figure 6

THIS PAGE BLANK (USPTO)

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
25 April 2002 (25.04.2002)

PCT

(10) International Publication Number
WO 02/033895 A3(51) International Patent Classification⁷: **H04L 12/26**,
12/56

(21) International Application Number: PCT/US01/32319

(22) International Filing Date: 17 October 2001 (17.10.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/241,450	17 October 2000 (17.10.2000)	US
60/275,206	12 March 2001 (12.03.2001)	US
09/903,423	10 July 2001 (10.07.2001)	US
09/903,441	10 July 2001 (10.07.2001)	US
09/923,924	6 August 2001 (06.08.2001)	US
09/960,623	20 September 2001 (20.09.2001)	US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US	60/241,450 (CIP)
Filed on	17 October 2000 (17.10.2000)
US	60/275,206 (CIP)
Filed on	12 March 2001 (12.03.2001)

US	09/903,441 (CIP)
Filed on	10 July 2001 (10.07.2001)
US	09/903,423 (CIP)
Filed on	10 July 2001 (10.07.2001)
US	09/923,924 (CIP)
Filed on	6 August 2001 (06.08.2001)
US	09/960,623 (CIP)
Filed on	20 September 2001 (20.09.2001)

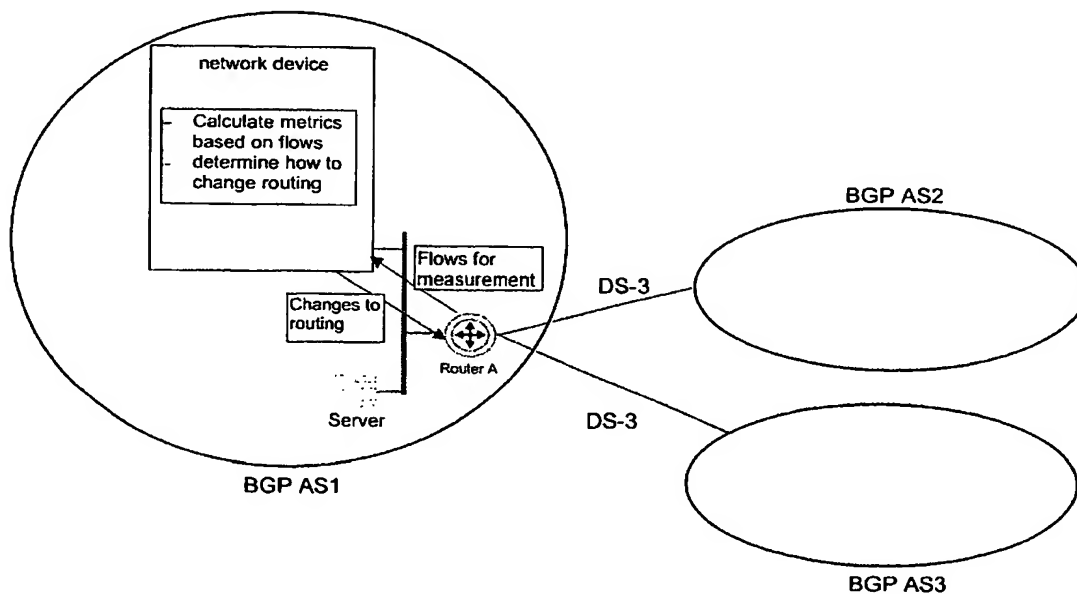
(71) Applicant (for all designated States except US): **ROUTE-SCIENCE TECHNOLOGIES, INC.** [US/US]; 167 2nd Avenue, San Mateo, CA 94401 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **LLOYD, Michael, A.** [US/US]; 160 Arundel Road, San Carlos, CA 94070 (US). **FINN, Sean, P.** [US/US]; 1533 Escondido Way, Belmont, CA 94002 (US). **KARAM, Mansour, J.** [LB/US]; 707 Continental Circle, #421, Mountain View, CA 94040 (US). **BALDONADO, Omar, C.** [US/US]; 700 Alester Avenue, Palo Alto, CA 94303 (US). **MADAN, Herbert, S.** [US/US]; 347 Blackfield Drive, Tiburon, CA 94920 (US). **MCGUIRE, James, G.** [US/US]; 2312 Gough Street, San Francisco, CA 94019 (US).

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR THE ASSESSMENT AND OPTIMIZATION OF NETWORK TRAFFIC



(57) Abstract: A system for the assessment of network performance criteria, and applying this criteria to the classification of network addresses into appropriate ranges, using these ranges to consolidate performance measurements for the associated addresses, and applying these metrics toward the optimization of the network towards performance or policy objectives.

WO 02/033895 A3



(74) **Agent:** SUZUE, Kenta; Wilson Sonsini Goodrich & Rosati, 650 Page Mill Road, Palo Alto, CA 94304-1050 (US).

(81) **Designated States (national):** AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian

patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

(88) **Date of publication of the international search report:**
7 August 2003

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

INTERNATIONAL SEARCH REPORT

Int. Application No
PCT/US 01/32319A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/26 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 00 38381 A (DIGITAL ISLAND INC) 29 June 2000 (2000-06-29)	80-86, 90, 91, 105, 107-115
Y	page 3, line 9 -page 5, line 5 claims --- -/--	92, 93, 99-101, 104, 106

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

9 April 2003

Date of mailing of the international search report

22 APR 2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

Int. Application No.

PCT/US 01/32319

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	PAPADOPOULOS C V ET AL: "PROTECTION AND ROUTING ALGORITHMS FOR NETWORK MANAGEMENT THE CASE OF TRANSMISSION NETWORKS" MICROPROCESSING AND MICROPROGRAMMING, - ELSEVIER SCIENCE PUBLISHERS, BV., AMSTERDAM, NL, vol. 38, no. 1 / 5, 1 September 1993 (1993-09-01), pages 163-170, XP000383771 ISSN: 0165-6074 sections 2.1 and 2.2	92,93
Y	US 5 884 047 A (AIKAWA HIDEYUKI ET AL) 16 March 1999 (1999-03-16) abstract claims	99-101, 104,106
A	US 5 467 345 A (RICETTA PETER ET AL) 14 November 1995 (1995-11-14) abstract claim 22	1,75,76, 78
A	US 5 787 253 A (MCCREERY TIMOTHY DAVID ET AL) 28 July 1998 (1998-07-28) column 2, line 50 - line 67 column 4, line 48 - line 62	1,75,76, 78
A	EP 0 504 537 A (IBM) 23 September 1992 (1992-09-23) claims	1,75,76, 78

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 01/32319

Box I Observations where certain claims were found unsearchable (Continuation of Item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of Item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☒ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☒ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. Claims: 1-79

Method and apparatus for the optimization of internetwork traffic

2. Claims: 80-115

A method for classifying network addresses into appropriated groups

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 01/32319

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0038381	A	29-06-2000	US 6130890 A AU 6140899 A TW 466853 B TW 435029 B WO 0038381 A1	10-10-2000 12-07-2000 01-12-2001 16-05-2001 29-06-2000
US 5884047	A	16-03-1999	JP 3361865 B2 JP 7162416 A GB 2284730 A ,B	07-01-2003 23-06-1995 14-06-1995
US 5467345	A	14-11-1995	FR 2720584 A1	01-12-1995
US 5787253	A	28-07-1998	NONE	
EP 0504537	A	23-09-1992	EP 0504537 A1 JP 4328926 A US 5271000 A	23-09-1992 17-11-1992 14-12-1993

Form PCT/ISA/210 (patent family annex) (July 1992)

THIS PAGE BLANK (USPTO)